

Intel® Architecture Instruction Set Extensions Programming Reference

319433-027

DECEMBER 2016

Intel technologies features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Intel does not guarantee the availability of these interfaces in any future product. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting <http://www.intel.com/design/literature.htm>.

Intel, the Intel logo, Intel Atom, Intel Core, Intel SpeedStep, MMX, Pentium, VTune, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 1997-2016, Intel Corporation. All Rights Reserved.

Revision History

Revision	Description	Date
-025	<ul style="list-style-type: none">Removed instructions that now reside in the Intel® 64 and IA-32 Architectures Software Developer's Manual.Minor updates to chapter 1.Updates to Table 2-1, Table 2-2 and Table 2-8 (leaf 07H) to indicate support for AVX512_4VNNIW and AVX512_4FMAPS.Minor update to Table 2-8 (leaf 15H) regarding ECX definition.Minor updates to Section 4.6.2 and Section 4.6.3 to clarify the effects of "suppress all exceptions".Footnote addition to CLWB instruction indicating operand encoding requirement.Removed PCOMMIT.	September 2016
-026	<ul style="list-style-type: none">Removed CLWB instruction; it now resides in the Intel® 64 and IA-32 Architectures Software Developer's Manual.Added additional 512-bit instruction extensions in chapter 6.	October 2016
-027	<ul style="list-style-type: none">Added TLB CPUID leaf in chapter 2.Added VPOPCNTD/Q instruction in chapter 6, and CPUID details in chapter 2.	December 2016

REVISION HISTORY

CHAPTER 1

FUTURE INTEL® ARCHITECTURE INSTRUCTION EXTENSIONS

1.1	About This Document.....	1-1
1.2	Intel® AVX-512 Instructions Architecture Overview.....	1-1
1.2.1	512-Bit Wide SIMD Register Support.....	1-2
1.2.2	32 SIMD Register Support.....	1-2
1.2.3	Eight Opmask Register Support.....	1-2
1.2.4	Instruction Syntax Enhancement.....	1-2
1.2.5	EVEX Instruction Encoding Support.....	1-3

CHAPTER 2

INTEL® AVX-512 APPLICATION PROGRAMMING MODEL

2.1	Detection of AVX-512 Foundation Instructions.....	2-1
2.2	Detection of 512-bit Instruction Groups of Intel® AVX-512 Family.....	2-2
2.3	Detection of Intel AVX-512 Instruction Groups Operating at 256 and 128-bit Vector Lengths.....	2-3
2.4	Accessing XMM, YMM AND ZMM Registers.....	2-4
2.5	Enhanced Vector Programming Environment Using EVEX Encoding.....	2-4
2.5.1	OPMASK Register to Predicate Vector Data Processing.....	2-5
2.5.1.1	Opmask Register K0.....	2-6
2.5.1.2	Example of Opmask Usages.....	2-6
2.5.2	OpMask Instructions.....	2-7
2.5.3	Broadcast.....	2-7
2.5.4	STATIC ROUNDING MODE AND SUPPRESS ALL EXCEPTIONS.....	2-8
2.5.5	Compressed Disp8*N Encoding.....	2-9
2.6	Memory Alignment.....	2-10
2.7	SIMD Floating-Point Exceptions.....	2-11
2.8	Instruction Exception Specification.....	2-11
2.9	CPUID Instruction.....	2-12
	CPUID—CPU Identification.....	2-12

CHAPTER 3

SYSTEM PROGRAMMING FOR INTEL® AVX-512

3.1	AVX-512 State, EVEX Prefix and Supported Operating Modes.....	3-1
3.2	AVX-512 State Management.....	3-1
3.2.1	Detection of ZMM and Opmask State Support.....	3-1
3.2.2	Enabling of ZMM and Opmask Register State.....	3-2
3.2.3	Enabling of SIMD Floating-Exception Support.....	3-3
3.2.4	The Layout of XSAVE State Save Area.....	3-3
3.2.5	XSAVE/XRSTOR Interaction with YMM State and MXCSR.....	3-5
3.2.6	XSAVE/XRSTOR/XSAVEOPT and Managing ZMM and Opmask States.....	3-6
3.3	Reset Behavior.....	3-7
3.4	Emulation.....	3-7
3.5	Writing floating-point exception handlers.....	3-7

CHAPTER 4

INTEL® AVX-512 INSTRUCTION ENCODING

4.1	Overview Section.....	4-1
4.2	Instruction Format and EVEX.....	4-1
4.3	Register Specifier Encoding and EVEX.....	4-3
4.3.1	Opmask Register Encoding.....	4-4
4.4	MASKING support in EVEX.....	4-4
4.5	Compressed displacement (disp8*N) support in EVEX.....	4-5
4.6	EVEX encoding of broadcast/Rounding/SAE Support.....	4-6
4.6.1	Embedded Broadcast Support in EVEX.....	4-6
4.6.2	Static Rounding Support in EVEX.....	4-6
4.6.3	SAE Support in EVEX.....	4-7

4.6.4	Vector Length Orthogonality	4-7
4.7	#UD equations for EVEX	4-7
4.7.1	State Dependent #UD	4-7
4.7.2	Opcode Independent #UD	4-8
4.7.3	Opcode Dependent #UD	4-8
4.8	Device Not Available	4-9
4.9	Scalar Instructions	4-10
4.10	Exception Classifications of EVEX-Encoded instructions	4-10
4.10.1	Exceptions Type E1 and E1NF of EVEX-Encoded Instructions	4-13
4.10.2	Exceptions Type E2 of EVEX-Encoded Instructions	4-15
4.10.3	Exceptions Type E3 and E3NF of EVEX-Encoded Instructions	4-16
4.10.4	Exceptions Type E4 and E4NF of EVEX-Encoded Instructions	4-18
4.10.5	Exceptions Type E5 and E5NF	4-20
4.10.6	Exceptions Type E6 and E6NF	4-22
4.10.7	Exceptions Type E7NM	4-24
4.10.8	Exceptions Type E9 and E9NF	4-25
4.10.9	Exceptions Type E10	4-27
4.10.10	Exception Type E11 (EVEX-only, mem arg no AC, floating-point exceptions)	4-29
4.10.11	Exception Type E12 and E12NP (VSIB mem arg, no AC, no floating-point exceptions)	4-30
4.11	Exception Classifications of Opmask instructions	4-32

CHAPTER 5

INSTRUCTION SET REFERENCE, A-Z

5.1	Interpreting Instruction Reference Pages	5-1
5.1.1	Instruction Format	5-1
	ADDPS—Add Packed Single-Precision Floating-Point Values (THIS IS AN EXAMPLE)	5-1
5.1.2	Opcode Column in the Instruction Summary Table	5-1
5.1.3	Instruction Column in the Instruction Summary Table	5-4
5.1.4	64/32 bit Mode Support column in the Instruction Summary Table	5-5
5.1.5	CPUID Support column in the Instruction Summary Table	5-5
5.1.5.1	Operand Encoding Column in the Instruction Summary Table	5-5
5.2	Summary of Terms	5-6
5.3	Ternary Bit Vector Logic Table	5-6
5.4	Instruction SET Reference	5-8
	VPERMI2B—Full Permute of Bytes from Two Tables Overwriting the Index	5-9
	VPERMT2B—Full Permute of Bytes from Two Tables Overwriting a Table	5-11
	See Exceptions Type E4NF.nb	5-12
	VPERMT2W/D/Q/PS/PD—Full Permute from Two Tables Overwriting one Table	5-13
	VPMADD52LUQ—Packed Multiply of Unsigned 52-bit Integers and Add the Low 52-bit Products to Qword Accumulators	5-18
	VPMADD52HUQ—Packed Multiply of Unsigned 52-bit Unsigned Integers and Add High 52-bit Products to 64-bit Accumulators	5-20
	VPMULTISHIFTQB – Select Packed Unaligned Bytes from Quadword Sources	5-22

CHAPTER 6

ADDITIONAL 512-BIT INSTRUCTION EXTENSIONS

6.1	INSTRUCTION Set Reference	6-1
	V4FMADDPS/V4FNMADDPS — Packed Single-Precision Floating-Point Fused Multiply-Add (4-iterations)	6-1
	V4FMADDSS/V4FNMADDSS — Scalar Single-Precision Floating-Point Fused Multiply-Add (4-iterations)	6-3
	VP4DPWSSD — Dot Product of Signed Words with Dword Accumulation (4-iterations)	6-5
	VP4DPWSSDS — Dot Product of Signed Words with Dword Accumulation and Saturation (4-iterations)	6-7
	VPOPCNTD/VPOPCNTQ — Return the Count of Number of Bits Set to 1 in DWORD/QWORD	6-9

TABLES

	PAGE
2-1	512-bit Instruction Groups in the Intel AVX-512 Family 2-2
2-2	Feature flag Collection Required of 256/128 Bit Vector Lengths for Each Instruction Group 2-4
2-3	Instruction Mnemonics That Do Not Support EVEX.128 Encoding 2-4
2-4	Characteristics of Three Rounding Control Interfaces 2-8
2-5	Static Rounding Mode 2-9
2-6	SIMD Instructions Requiring Explicitly Aligned Memory 2-10
2-7	Instructions Not Requiring Explicit Memory Alignment 2-11
2-8	Information Returned by CPUID Instruction 2-13
2-9	Highest CPUID Source Operand for Intel 64 and IA-32 Processors 2-25
2-10	Processor Type Field 2-26
2-11	Feature Information Returned in the ECX Register 2-28
2-12	More on Feature Information Returned in the EDX Register 2-29
2-13	Encoding of Cache and TLB Descriptors 2-31
2-14	Structured Extended Feature Leaf, Function 0, EBX Register 2-34
2-15	Processor Brand String Returned with Pentium 4 Processor 2-38
2-16	Mapping of Brand Indices; and Intel 64 and IA-32 Processor Brand Strings 2-40
3-1	XCRO Processor State Components 3-2
3-2	CR4 Bits for AVX-512 Foundation Instructions Technology Support 3-3
3-3	Layout of XSAVE Area For Processor Supporting YMM State 3-4
3-4	XSAVE Header Format 3-4
3-5	XSAVE Save Area Layout for YMM_Hi128 State (Ext_Save_Area_2) 3-4
3-6	XSAVE Save Area Layout for Opmask Registers 3-5
3-7	XSAVE Save Area Layout for ZMM State of the High 256 Bits of ZMM0-ZMM15 Registers 3-5
3-8	XSAVE Save Area Layout for ZMM State of ZMM16-ZMM31 Registers 3-5
3-9	XRSTOR Action on MXCSR, XMM Registers, YMM Registers 3-6
3-10	XSAVE Action on MXCSR, XMM, YMM Register 3-6
3-11	Processor Supplied Init Values XRSTOR May Use 3-7
4-1	EVEX Prefix Bit Field Functional Grouping 4-2
4-2	32-Register Support in 64-bit Mode Using EVEX with Embedded REX Bits 4-3
4-3	EVEX Encoding Register Specifiers in 32-bit Mode 4-4
4-4	Opmask Register Specifier Encoding 4-4
4-5	Compressed Displacement (DISP8*N) Affected by Embedded Broadcast 4-5
4-6	EVEX DISP8*N For Instructions Not Affected by Embedded Broadcast 4-6
4-7	EVEX Embedded Broadcast/Rounding/SAE and Vector Length on Vector Instructions 4-7
4-8	OS XSAVE Enabling Requirements of Instruction Categories 4-8
4-9	Opcode Independent, State Dependent EVEX Bit Fields 4-8
4-10	#UD Conditions of Operand-Encoding EVEX Prefix Bit Fields 4-8
4-11	#UD Conditions of Opmask Related Encoding Field 4-9
4-12	#UD Conditions Dependent on EVEX.b Context 4-9
4-13	EVEX-Encoded Instruction Exception Class Summary 4-10
4-14	EVEX Instructions in each Exception Class 4-11
4-15	Type E1 Class Exception Conditions 4-13
4-16	Type E1NF Class Exception Conditions 4-14
4-17	Type E2 Class Exception Conditions 4-15
4-18	Type E3 Class Exception Conditions 4-16
4-19	Type E3NF Class Exception Conditions 4-17
4-20	Type E4 Class Exception Conditions 4-18
4-21	Type E4NF Class Exception Conditions 4-19
4-22	Type E5 Class Exception Conditions 4-20
4-23	Type E5NF Class Exception Conditions 4-21
4-24	Type E6 Class Exception Conditions 4-22
4-25	Type E6NF Class Exception Conditions 4-23
4-26	Type E7NM Class Exception Conditions 4-24
4-27	Type E9 Class Exception Conditions 4-25
4-28	Type E9NF Class Exception Conditions 4-26

4-29	Type E10 Class Exception Conditions	4-27
4-30	Type E10NF Class Exception Conditions	4-28
4-31	Type E11 Class Exception Conditions	4-29
4-32	Type E12 Class Exception Conditions	4-30
4-33	Type E12NP Class Exception Conditions	4-31
4-34	TYPE K20 Exception Definition (VEX-Encoded OpMask Instructions w/o Memory Arg)	4-32
4-35	TYPE K21 Exception Definition (VEX-Encoded OpMask Instructions Addressing Memory)	4-33
5-1	Low 8 columns of the 16x16 Map of VPTERNLOG Boolean Logic Operations.....	5-7
5-2	Low 8 columns of the 16x16 Map of VPTERNLOG Boolean Logic Operations.....	5-8

FIGURES

	PAGE
Figure 1-1. 512-Bit Wide Vectors and SIMD Register Set.....	1-2
Figure 2-1. Procedural Flow of Application Detection of AVX-512 Foundation Instructions.....	2-1
Figure 2-2. Procedural Flow of Application Detection of 512-bit Instruction Groups	2-2
Figure 2-3. Procedural Flow of Application Detection of Intel AVX-512 Instructions Operating at Vector Lengths < 512.....	2-3
Figure 2-4. Version Information Returned by CPUID in EAX.....	2-26
Figure 2-5. Feature Information Returned in the ECX Register	2-27
Figure 2-6. Feature Information Returned in the EDX Register	2-29
Figure 2-7. Determination of Support for the Processor Brand String.....	2-37
Figure 2-8. Algorithm for Extracting Maximum Processor Frequency	2-39
Figure 3-1. Bit Vector and XCRO Layout of Extended Processor State Components.....	3-2
Figure 4-1. AVX-512 Instruction Format and the EVEX Prefix.....	4-1
Figure 4-2. Bit Field Layout of the EVEX Prefix.....	4-2
Figure 6-1. Register Source-Block Dot Product of Two Signed Word Operands with Doubleword Accumulation.....	6-5

CHAPTER 1

FUTURE INTEL® ARCHITECTURE INSTRUCTION EXTENSIONS

1.1 ABOUT THIS DOCUMENT

This document describes the software programming interfaces of Intel® architecture instruction extensions which may be included in future Intel processor generations. Intel does not guarantee the availability of these interfaces in any future product.

The instruction set extensions cover a diverse range of application domains and programming usages. The 512-bit SIMD vector SIMD extensions, referred to as Intel® Advanced Vector Extensions 512 (Intel® AVX-512) instructions, deliver comprehensive set of functionality and higher performance than Intel® AVX and Intel® AVX2 instructions. Intel AVX, Intel AVX2 and many Intel AVX-512 instructions are covered in *Intel® 64 and IA-32 Architectures Software Developer's Manual sets*. The reader can refer to them for basic and more background information related to various features referenced in this document.

The base of the 512-bit SIMD instruction extensions are referred to as Intel AVX-512 Foundation instructions. They include extensions of the AVX and AVX2 family of SIMD instructions but are encoded using a new encoding scheme with support for 512-bit vector registers, up to 32 vector registers in 64-bit mode, and conditional processing using opmask registers.

Chapters 2 through 5 are devoted to the programming interfaces of the AVX-512 Foundation instruction set, additional 512-bit instruction extensions in the Intel AVX-512 family targeting broad application domains, and instruction set extensions encoded using the EVEX prefix encoding scheme to operate at vector lengths smaller than 512-bits.

Chapter 6 describes instruction set extensions that offer software tools with capability to address memory protection issues such as buffer overruns.

1.2 INTEL® AVX-512 INSTRUCTIONS ARCHITECTURE OVERVIEW

Intel AVX-512 Foundation instructions are a natural extension to AVX and AVX2. It introduces the following architectural enhancements:

- Support for 512-bit wide vectors and SIMD register set. 512-bit register state is managed by the operating system using XSAVE/XRSTOR instructions introduced in 45 nm Intel 64 processors (see *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2C*, and *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*).
- Support for 16 new, 512-bit SIMD registers (for a total of 32 SIMD registers, ZMM0 through ZMM31) in 64-bit mode. The extra 16 registers state is managed by the operating system using XSAVE/XRSTOR/XSAVEOPT.
- Support for 8 new opmask registers (k0 through k7) used for conditional execution and efficient merging of destination operands. Again, the opmask register state is managed by the operating system using XSAVE/XRSTOR/XSAVEOPT instructions
- A new encoding prefix (referred to as EVEX) to support additional vector length encoding up to 512 bits. The EVEX prefix builds upon the foundations of VEX prefix, to provide compact, efficient encoding for functionality available to VEX encoding plus the following enhanced vector capabilities:
 - opmasks
 - embedded broadcast
 - instruction prefix-embedded rounding control
 - compressed address displacements

1.2.1 512-Bit Wide SIMD Register Support

AVX-512 instructions support 512-bit wide SIMD registers (ZMM0-ZMM31). The lower 256-bits of the ZMM registers are aliased to the respective 256-bit YMM registers and the lower 128-bit are aliased to the respective 128-bit XMM registers.

1.2.2 32 SIMD Register Support

AVX-512 instructions also support for 32 SIMD registers in 64-bit mode (XMM0-XMM31, YMM0-YMM31 and ZMM0-ZMM31). The number of available vector registers in 32-bit mode is still 8.

1.2.3 Eight Opmask Register Support

AVX-512 instructions support 8 opmask registers (k0-k7). The width of each opmask register is architecturally defined of size MAX_KL (64 bits). Seven of the eight opmask registers (k1-k7) can be used in conjunction with EVEX-encoded AVX-512 Foundation instructions to provide conditional execution and efficient merging of data elements in the destination operand. The encoding of opmask register k0 is typically used when all data elements (unconditional processing) are desired. Additionally, the opmask registers are also used as vector flags/element-level vector sources to introduce novel SIMD functionality as seen in new instructions such as VCOMPRESSPS.




Figure 1-1. 512-Bit Wide Vectors and SIMD Register Set

1.2.4 Instruction Syntax Enhancement

The architecture of EVEX encoding enhances vector instruction encoding scheme in the following way:

- 512-bit vector-length, up to 32 ZMM registers, and enhanced vector programming environment are supported using the enhanced VEX (EVEX).

The EVEX prefix provides more encodable bit fields than VEX prefix. In addition to encoding 32 ZMM registers in 64-bit mode, instruction encoding using the EVEX can directly encode 7 (out of 8) opmask register operands to provide conditional processing in vector instruction programming. The enhanced vector programming environment can be explicitly expressed in the instruction syntax to include the following elements:

- An opmask operand: the opmask registers are expressed using the notation “k1” through “k7”. An EVEX-encoded instruction supporting conditional vector operation using the opmask register k1 is expressed by attaching the notation {k1} next to the destination operand. The use of this feature is optional for most instructions. There are two types of masking (merging and zeroing) differentiated using the EVEX.z bit ({z} in instruction signature).

- Embedded broadcast may be supported for some instructions on the source operand that can be encoded as a memory vector. Data elements of a memory vector may be conditionally fetched or written to.
- For instruction syntax that operates only on floating-point data in SIMD registers with rounding semantics, the EVEX can provide explicit rounding control within the EVEX bit fields at either scalar or 512-bit vector length.

In AVX-512 instructions, vector addition of all elements of the source operands can be expressed in the same syntax as AVX instruction:

```
VADDPS zmm1, zmm2, zmm3
```

Additionally, the EVEX encoding scheme of AVX-512 Foundation can express conditional vector addition as

```
VADDPS zmm1 {k1}{z}, zmm2, zmm3
```

where

- conditional processing and updates to destination is expressed with an opmask register,
- zeroing behavior of the opmask selected destination element is expressed by the {z} modifier (with merging as the default if no modifier specified),

Note that some SIMD instructions supporting three-operand syntax but processing only less or equal than 128-bits of data are considered part of the 512-bit SIMD instruction set extensions, because bits MAX_VL-1:128 of the destination register are zeroed by the processor. The same rule applies to instructions operating on 256-bits of data where bits MAX_VL-1:256 of the destination register are zeroed.

1.2.5 EVEX Instruction Encoding Support

Intel AVX-512 instructions employ a new encoding prefix, referred to as EVEX, in the Intel 64 and IA-32 instruction encoding format. Instruction encoding using the EVEX prefix provides the following capabilities:

- Direct encoding of a SIMD register operand within EVEX (similar to VEX). This provides instruction syntax support for three source operands.
- Compaction of REX prefix functionality and extended SIMD register encoding: The equivalent REX-prefix compaction functionality offered by the VEX prefix is provided within EVEX. Furthermore, EVEX extends the operand encoding capability to allow direct addressing of up to 32 ZMM registers in 64-bit mode.
- Compaction of SIMD prefix functionality and escape byte encoding: The functionality of SIMD prefix (66H, F2H, F3H) on opcode is equivalent to an opcode extension field to introduce new processing primitives. This functionality is provided in the VEX prefix encoding scheme and employed within the EVEX prefix. Similarly, the functionality of the escape opcode byte (0FH) and two-byte escape (0F38H, 0F3AH) are also compacted within the EVEX prefix encoding.
- Most EVEX-encoded SIMD numeric and data processing instruction semantics with memory operand have relaxed memory alignment requirements than instructions encoded using SIMD prefixes (see Section 2.6, “Memory Alignment”).
- Direct encoding of a opmask operand within the EVEX prefix. This provides instruction syntax support for conditional vector-element operation and merging of destination operand using an opmask register (k1-k7).
- Direct encoding of a broadcast attribute for instructions with a memory operand source. This provides instruction syntax support for elements broadcasting of the second operand before being used in the actual operation.
- Compressed memory address displacements for a more compact instruction encoding byte sequence.

EVEX encoding applies to SIMD instructions operating on XMM, YMM and ZMM registers. EVEX is not supported for instructions operating on MMX or x87 registers. Details of EVEX instruction encoding are discussed in Chapter 4.

CHAPTER 2

INTEL® AVX-512 APPLICATION PROGRAMMING MODEL

The application programming model for AVX-512 Foundation instructions and several member groups of the Intel® AVX-512 family (described in Chapter 5) extend from that of Intel AVX and Intel AVX2 with differences detailed in this chapter.

2.1 DETECTION OF AVX-512 FOUNDATION INSTRUCTIONS

The majority of AVX-512 Foundation instructions are encoded using the EVEX encoding scheme. EVEX-encoded instructions can operate on the 512-bit ZMM register state plus 8 opmask registers. The opmask instructions in AVX-512 Foundation instructions operate only on opmask registers or with a general purpose register. System software requirements to support ZMM state and opmask instructions are described in Chapter 3, “System Programming For Intel® AVX-512”.

Processor support of AVX-512 Foundation instructions is indicated by CPUID. (EAX=07H, ECX=0):EBX.AVX512F[bit 16] = 1. Detection of AVX-512 Foundation instructions operating on ZMM states and opmask registers need to follow the general procedural flow in Figure 2-1.

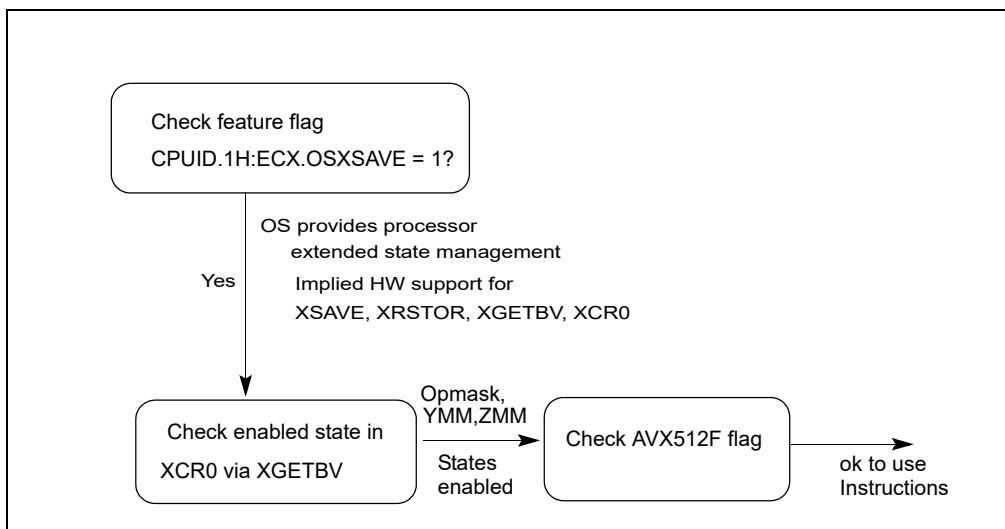


Figure 2-1. Procedural Flow of Application Detection of AVX-512 Foundation Instructions

Prior to using AVX-512 Foundation instructions, the application must identify that the operating system supports the XGETBV instruction, the ZMM register state, in addition to processor's support for ZMM state management using XSAVE/XRSTOR and AVX-512 Foundation instructions. The following simplified sequence accomplishes both and is strongly recommended.

- 1) Detect CPUID.1:ECX.OSXSAVE[bit 27] = 1 (XGETBV enabled for application use¹).
- 2) Execute XGETBV and verify that XCR0[7:5] = '111b' (OPMASK state, upper 256-bit of ZMM0-ZMM15 and ZMM16-ZMM31 state are enabled by OS) and that XCR0[2:1] = '11b' (XMM state and YMM state are enabled by OS).
- 3) Detect CPUID.0x7.0:EBX.AVX512F[bit 16] = 1.

1. If CPUID.01H:ECX.OSXSAVE reports 1, it also indirectly implies the processor supports XSAVE, XRSTOR, XGETBV, processor extended state bit vector XCR0 register. Thus an application may streamline the checking of CPUID feature flags for XSAVE and OSXSAVE. XSETBV is a privileged instruction.

2.2 DETECTION OF 512-BIT INSTRUCTION GROUPS OF INTEL® AVX-512 FAMILY

In addition to the Intel AVX-512 Foundation instructions, Intel AVX-512 family provides several additional 512-bit extensions in groups of instructions, each group is enumerated by a CPUID leaf 7 feature flag and can be encoded via EVEX.L'L field to support operation at vector lengths smaller than 512 bits. These instruction groups are listed in Table 2-1.

Table 2-1. 512-bit Instruction Groups in the Intel AVX-512 Family

CPUID Leaf 7 Feature Flag Bit	Feature Flag abbreviation of 512-bit Instruction Group	SW Detection Flow
CPUID.(EAX=07H, ECX=0):EBX[bit 16]	AVX512F (AVX-512 Foundation)	Figure 2-1
CPUID.(EAX=07H, ECX=0):EBX[bit 28]	AVX512CD	Figure 2-2
CPUID.(EAX=07H, ECX=0):EBX[bit 17]	AVX512DQ	Figure 2-2
CPUID.(EAX=07H, ECX=0):EBX[bit 30]	AVX512BW	Figure 2-2
CPUID.(EAX=07H, ECX=0):EBX[bit 21]	AVX512IFMA	Figure 2-2
CPUID.(EAX=07H, ECX=0):ECX[bit 01]	AVX512VBMI	Figure 2-2
CPUID.(EAX=07H, ECX=0):ECX[bit 14]	AVX512_VPOPCNTDQ	Figure 2-2
CPUID.(EAX=07H, ECX=0):EDX[bit 02]	AVX512_4VNNIW	Figure 2-2
CPUID.(EAX=07H, ECX=0):EDX[bit 03]	AVX512_4FMAPS	Figure 2-2

Software must follow the detection procedure for the 512-bit AVX-512 Foundation instructions as described in Section 2.1.

Detection of other 512-bit sibling instruction groups listed in Table 2-1 (excluding AVX512F) follows the procedure described in Figure 2-2:

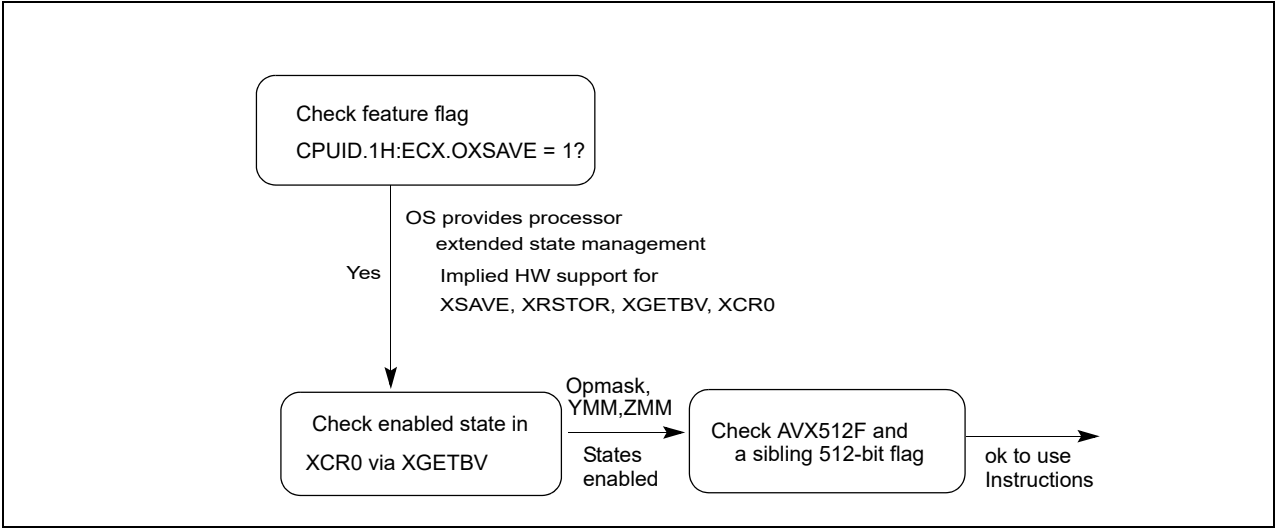


Figure 2-2. Procedural Flow of Application Detection of 512-bit Instruction Groups

To illustrate the detection procedure for 512-bit instructions enumerated by AVX512CD, the following sequence is strongly recommended.

- 1) Detect CPUID.1:ECX.OSXSAVE[bit 27] = 1 (XGETBV enabled for application use).

2) Execute XGETBV and verify that $\text{XCR0}[7:5] = '111b'$ (OPMASK state, upper 256-bit of ZMM0-ZMM15 and ZMM16-ZMM31 state are enabled by OS) and that $\text{XCR0}[2:1] = '11b'$ (XMM state and YMM state are enabled by OS).

3) Verify both $\text{CPUID.0x7.0:EBX.AVX512F}[\text{bit } 16] = 1$, $\text{CPUID.0x7.0:EBX.AVX512CD}[\text{bit } 28] = 1$.

Similarly, the detection procedure for enumerating 512-bit instructions reported by AVX512DW follows the same flow.

2.3 DETECTION OF INTEL AVX-512 INSTRUCTION GROUPS OPERATING AT 256 AND 128-BIT VECTOR LENGTHS

For each of the 512-bit instruction groups in the Intel AVX-512 family listed in Table 2-1, EVEX encoding scheme may support a vast majority of these instructions operating at 256-bit or 128-bit (if applicable) vector lengths. This encoding support for vector lengths smaller than 512-bits is indicated by $\text{CPUID.(EAX=07H, ECX=0):EBX}[\text{bit } 31]$, abbreviated as AVX512VL.

The AVX512VL flag alone is never sufficient to determine a given Intel AVX-512 instruction may be encoded at vector lengths smaller than 512 bits. Software must use the procedure described in Figure 2-3 and Table 2-2:




Figure 2-3. Procedural Flow of Application Detection of Intel AVX-512 Instructions Operating at Vector Lengths < 512

To illustrate the procedure described in Figure 2-3 and Table 2-2 for software to use EVEX.256 encoded VPCONFLICT, the following sequence is strongly recommended.

1) Detect $\text{CPUID.1:ECX.OSXSAVE}[\text{bit } 27] = 1$ (XGETBV enabled for application use)

2) Execute XGETBV and verify that $\text{XCR0}[7:5] = '111b'$ (OPMASK state, upper 256-bit of ZMM0-ZMM15 and ZMM16-ZMM31 state are enabled by OS) and that $\text{XCR0}[2:1] = '11b'$ (XMM state and YMM state are enabled by OS).

3) Verify $\text{CPUID.0x7.0:EBX.AVX512F}[\text{bit } 16] = 1$, $\text{CPUID.0x7.0:EBX.AVX512CD}[\text{bit } 28] = 1$, and $\text{CPUID.0x7.0:EBX.AVX512VL}[\text{bit } 31] = 1$.

Table 2-2. Feature flag Collection Required of 256/128 Bit Vector Lengths for Each Instruction Group

Usage of 256/128 Vector Lengths	Feature Flag Collection to Verify
AVX512F	AVX512F & AVX512VL
AVX512CD	AVX512F & AVX512CD & AVX512VL
AVX512DQ	AVX512F & AVX512DQ & AVX512VL
AVX512BW	AVX512F & AVX512BW & AVX512VL
AVX512IFMA	AVX512F & AVX512IFMA & AVX512VL
AVX512VBMI	AVX512F & AVX512VBMI & AVX512VL
AVX512_4FMAPS	AVX512F & AVX512_4FMAPS & AVX512VL
AVX512_4VNNIW	AVX512F & AVX512_4VNNIW & AVX512VL
AVX512_VPOPCNTDQ	AVX512F & AVX512_VPOPCNTDQ & AVX512VL

In some specific cases, AVX512VL may only support EVEX.256 encoding but not EVEX.128. These are listed in Table 2-3.

Table 2-3. Instruction Mnemonics That Do Not Support EVEX.128 Encoding

Instruction Group	Instruction Mnemonics Supporting EVEX.256 Only Using AVX512VL
AVX512F	VBROADCASTSD, VBROADCASTF32X4, VEXTRACTI32X4, VINSERTF32X4, VINSERTI32X4, VPERMD, VPERMPD, VPERMPS, VPERMQ, VSHUFF32X4, VSHUFF64X2, VSHUFI32X4, VSHUFI64X2
AVX512CD	
AVX512DQ	VBROADCASTF32X2, VBROADCASTF64X2, VBROADCASTI32X4, VBROADCASTI64X2, VEXTRACTI64X2, VINSERTF64X2, VINSERTI64X2,
AVX512BW	

2.4 ACCESSING XMM, YMM AND ZMM REGISTERS

The lower 128 bits of a YMM register is aliased to the corresponding XMM register. Legacy SSE instructions (i.e. SIMD instructions operating on XMM state but not using the VEX prefix, also referred to non-VEX encoded SIMD instructions) will not access the upper bits (MAX_VL-1:128) of the YMM registers. AVX and FMA instructions with a VEX prefix and vector length of 128-bits zeroes the upper 128 bits of the YMM register.

Upper bits of YMM registers (255:128) can be read and written by many instructions with a VEX.256 prefix.

XSAVE and XRESTOR may be used to save and restore the upper bits of the YMM registers.

The lower 256 bits of a ZMM register are aliased to the corresponding YMM register. Legacy SSE instructions (i.e. SIMD instructions operating on XMM state but not using the VEX prefix, also referred to non-VEX encoded SIMD instructions) will not access the upper bits (MAX_VL-1:128) of the ZMM registers, where MAX_VL is maximum vector length (currently 512 bits). AVX and FMA instructions with a VEX prefix and vector length of 128-bits zero the upper 384 bits of the ZMM register, while VEX prefix and vector length of 256-bits zeros the upper 256 bits of the ZMM register.

Upper bits of ZMM registers (511:256) can be read and written by instructions with an EVEX.512 prefix.

2.5 ENHANCED VECTOR PROGRAMMING ENVIRONMENT USING EVEX ENCODING

EVEX-encoded AVX-512 instructions support an enhanced vector programming environment. The enhanced vector programming environment uses the combination of EVEX bit-field encodings and a set of eight opmask registers to provide the following capabilities:

- Conditional vector processing of EVEX-encoded instruction. Opmask registers k1 through k7 can be used to conditionally govern the per-data-element computational operation and the per-element updates to the destination operand of an AVX-512 Foundation instruction. Each bit of the opmask register governs one vector element operation (a vector element can be of 32 bits or 64 bits).
- In addition to providing predication control on vector instructions via EVEX bit-field encoding, the opmask registers can also be used similarly to general-purpose registers as source/destination operands using modR/M encoding for non-mask-related instructions. In this case, an opmask register k0 through k7 can be selected.
- In 64-bit mode, 32 vector registers can be encoded using EVEX prefix.
- Broadcast may be supported for some instructions on the operand that can be encoded as a memory vector. The data elements of a memory vector may be conditionally fetched or written to, and the vector size is dependent on the data transformation function.
- Flexible rounding control for register-to-register flavor of EVEX encoded 512-bit and scalar instructions. Four rounding modes are supported by direct encoding within the EVEX prefix overriding MXCSR settings.
- Broadcast of one element to the rest of the destination vector register.
- Compressed 8-bit displacement encoding scheme to increase the instruction encoding density for instructions that normally require disp32 syntax.

2.5.1 OPMASK Register to Predicate Vector Data Processing

AVX-512 instructions using EVEX encodes a predicate operand to conditionally control per-element computational operation and updating of result to the destination operand. The predicate operand is known as the opmask register. The opmask is a set of eight architectural registers of size MAX_KL (64-bit). Note that from this set of 8 architectural registers, only k1 through k7 can be addressed as predicate operand. k0 can be used as a regular source or destination but cannot be encoded as a predicate operand. Note also that a predicate operand can be used to enable memory fault-suppression for some instructions with a memory operand (source or destination).

As a predicate operand, the opmask registers contain one bit to govern the operation/update to each data element of a vector register. In general, opmask registers can support instructions with element sizes: single-precision floating-point (float32), integer doubleword (int32), double-precision floating-point (float64), integer quadword (int64). The length of a opmask register, MAX_KL, is sufficient to handle up to 64 elements with one bit per element, i.e. 64 bits. Masking is supported in most of the AVX-512 instructions. For a given vector length, each instruction accesses only the number of least significant mask bits that are needed based on its data type. For example, AVX-512 Foundation instructions operating on 64-bit data elements with a 512-bit vector length, only use the 8 least significant bits of the opmask register.

An opmask register affects an AVX-512 instruction at per-element granularity. So, any numeric or non-numeric operation of each data element and per-element updates of intermediate results to the destination operand are predicated on the corresponding bit of the opmask register.

An opmask serving as a predicate operand in AVX-512 obeys the following properties:

- The instruction's operation is not performed for an element if the corresponding opmask bit is not set. This implies that no exception or violation can be caused by an operation on a masked-off element. Consequently, no MXCSR exception flag is updated as a result of a masked-off operation.
- A destination element is not updated with the result of the operation if the corresponding writemask bit is not set. Instead, the destination element value must be preserved (merging-masking) or it must be zeroed out (zeroing-masking).
- For some instructions with a memory operand, memory faults are suppressed for elements with a mask bit of 0.

Note that this feature provides a versatile construct to implement control-flow predication as the mask in effect provides a merging behavior for AVX-512 vector register destinations. As an alternative the masking can be used for zeroing instead of merging, so that the masked out elements are updated with 0 instead of preserving the old value. The zeroing behavior is provided to remove the implicit dependency on the old value when it is not needed.

Most instructions with masking enabled accept both forms of masking. Instructions that must have EVEX.aaa bits different than 0 (gather and scatter) and instructions that write to memory only accept merging-masking.

It's important to note that the per-element destination update rule also applies when the destination operand is a memory location. Vectors are written on a per element basis, based on the opmask register used as a predicate operand.

The value of an opmask register can be:

- generated as a result of a vector instruction (e.g. CMP)
- loaded from memory
- loaded from GPR register
- or modified by mask-to-mask operations

Opmask registers can be used for purposes outside of predication. For example, they can be used to manipulate sparse sets of elements from a vector or used to set the EFLAGS based on the 0/0xFFFFFFFF/other status of the OR of two opmask registers.

2.5.1.1 Opmask Register K0

The only exception to the opmask rules described above is that opmask k0 can not be used as a predicate operand. Opmask k0 cannot be encoded as a predicate operand for a vector operation; the encoding value that would select opmask k0 will instead selects an implicit opmask value of 0xFFFFFFFF, thereby effectively disabling masking. Opmask register k0 can still be used for any instruction that takes opmask register(s) as operand(s) (either source or destination).

Note that certain instructions implicitly use the opmask as an extra destination operand. In such cases, trying to use the “no mask” feature will translate into a #UD fault being raised.

2.5.1.2 Example of Opmask Usages

The example below illustrates predicated vector add operation and predicated updates of added results into the destination operand. The initial state of vector registers zmm0, zmm1, and zmm2 and k3 are:

MSB.....LSB

```
zmm0 =
[ 0x00000003 0x00000002 0x00000001 0x00000000 ] (bytes 15 through 0)
[ 0x00000007 0x00000006 0x00000005 0x00000004 ] (bytes 31 through 16)
[ 0x0000000B 0x0000000A 0x00000009 0x00000008 ] (bytes 47 through 32)
[ 0x0000000F 0x0000000E 0x0000000D 0x0000000C ] (bytes 63 through 48)
```

```
zmm1 =
[ 0x0000000F 0x0000000F 0x0000000F 0x0000000F ] (bytes 15 through 0)
[ 0x0000000F 0x0000000F 0x0000000F 0x0000000F ] (bytes 31 through 16)
[ 0x0000000F 0x0000000F 0x0000000F 0x0000000F ] (bytes 47 through 32)
[ 0x0000000F 0x0000000F 0x0000000F 0x0000000F ] (bytes 63 through 48)
```

```
zmm2 =
[ 0xAAAAAAAA 0xAAAAAAAA 0xAAAAAAAA 0xAAAAAAAA ] (bytes 15 through 0)
[ 0BBBBBBBB 0BBBBBBBB 0BBBBBBBB 0BBBBBBBB ] (bytes 31 through 16)
[ 0xCCCCCCCC 0xCCCCCCCC 0xCCCCCCCC 0xCCCCCCCC ] (bytes 47 through 32)
[ 0xDDDDDDDD 0xDDDDDDDD 0xDDDDDDDD 0xDDDDDDDD ] (bytes 63 through 48)
```

```
k3 = 0x8F03 (1000 1111 0000 0011)
```

An opmask register serving as a predicate operand is expressed as a curly-braces-enclosed decorator following the first operand in the Intel assembly syntax. Given this state, we will execute the following instruction:

```
vpaddd zmm2 {k3}, zmm0, zmm1
```

The vpaddd instruction performs 32-bit integer additions on each data element conditionally based on the corresponding bit value in the predicate operand k3. Since per-element operations are not operated if the corresponding bit of the predicate mask is not set, the intermediate result is:

```
[ ***** ***** 0x00000010 0x0000000F ] (bytes 15 through 0)
[ ***** ***** ***** ***** ] (bytes 31 through 16)
[ 0x0000001A 0x00000019 0x00000018 0x00000017 ] (bytes 47 through 32)
[ 0x0000001E ***** ***** ***** ] (bytes 63 through 48)
```

where "*****" indicates that no operation is performed.

This intermediate result is then written into the destination vector register, zmm2, using the opmask register k3 as the writemask, producing the following final result:

```
zmm2 =
[ 0xAAAAAAAA 0xAAAAAAAA 0x00000010 0x0000000F ] (bytes 15 through 0)
[ 0BBBBBBBBB 0BBBBBBBBB 0BBBBBBBBB 0BBBBBBBBB ] (bytes 31 through 16)
[ 0x0000001A 0x00000019 0x00000018 0x00000017 ] (bytes 47 through 32)
[ 0x0000001E 0xDDDDDDDD 0xDDDDDDDD 0xDDDDDDDD ] (bytes 63 through 48)
```

Note that for a 64-bit instruction (say vaddpd), only the 8 LSB of mask k3 (0x03) would be used to identify the predicate operation on each one of the 8 elements of the source/destination vectors.

2.5.2 OpMask Instructions

AVX-512 Foundation instructions provide a collection of opmask instructions that allow programmers to set, copy, or operate on the contents of a given opmask register. There are three types of opmask instructions:

- **Mask read/write instructions:** These instructions move data between a general-purpose integer register or memory and an opmask mask register, or between two opmask registers. For example:
 - kmovw k1, ebx; move lower 16 bits of ebx to k1.
- **Flag instructions:** This category, consisting of instructions that modify EFLAGS based on the content of opmask registers.
 - kortestw k1, k2; OR registers k1 and k2 and updated EFLAGS accordingly.
- **Mask logical instructions:** These instructions perform standard bitwise logical operations between opmask registers.
 - kandw k1, k2, k3; AND lowest 16 bits of registers k2 and k3, leaving the result in k1.

2.5.3 Broadcast

EVEX encoding provides a bit-field to encode data broadcast for some load-op instructions, i.e. instructions that load data from memory and perform some computational or data movement operation. A source element from memory can be broadcasted (repeated) across all the elements of the effective source operand (up to 16 times for 32-bit data element, up to 8 times for 64-bit data element). This is useful when we want to reuse the same scalar

operand for all the operations in a vector instruction. Broadcast is only enabled on instructions with an element size of 32 bits or 64 bits. Byte and word instructions do not support embedded broadcast.

The functionality of data broadcast is expressed as a curly-braces-enclosed decorator following the last register/memory operand in the Intel assembly syntax.

For instance:

```
vmulps zmm1, zmm2, [rax] {1to16}
```

The {1to16} primitive loads one float32 (single precision) element from memory, replicates it 16 times to form a vector of 16 32-bit floating-point elements, multiplies the 16 float32 elements with the corresponding elements in the first source operand vector, and put each of the 16 results into the destination operand.

AVX-512 instructions with store semantics and pure load instructions do not support broadcast primitives.

```
vmovaps [rax] {k3}, zmm19
```

In contrast, the k3 opmask register is used as the predicate operand in the above example. Only the store operation on data elements corresponding to the non-zero bits in k3 will be performed.

2.5.4 STATIC ROUNDING MODE AND SUPPRESS ALL EXCEPTIONS

In previous SIMD instruction extensions, rounding control is generally specified in MXCSR, with a handful of instructions providing per-instruction rounding override via encoding fields within the imm8 operand. AVX-512 offers a more flexible encoding attribute to override MXCSR-based rounding control for floating-pointing instruction with rounding semantic. This rounding attribute embedded in the EVEX prefix is called Static (per instruction) Rounding Mode or Rounding Mode override. This attribute allows programmers to statically apply a specific arithmetic rounding mode irrespective of the value of RM bits in MXCSR. It is available only to register-to-register flavors of EVEX-encoded floating-point instructions with rounding semantic. The differences between these three rounding control interfaces are summarized in Table 2-4.

Table 2-4. Characteristics of Three Rounding Control Interfaces

Rounding Interface	Static Rounding Override	Imm8 Embedded Rounding Override	MXCSR Rounding Control
Semantic Requirement	FP rounding	FP rounding	FP rounding
Prefix Requirement	EVEX.B = 1	NA	NA
Rounding Control	EVEX.L'L	IMM8[1:0] or MXCSR.RC (depending on IMM8[2])	MXCSR.RC
Suppress All Exceptions (SAE)	Implied	no	no
SIMD FP Exception #XF	All suppressed	Can raise #I, #P (unless SPE is set)	MXCSR masking controls
MXCSR flag update	No	yes (except PE if SPE is set)	Yes
Precedence	Above MXCSR.RC	Above EVEX.L'L	Default
Scope	512-bit, reg-reg, Scalar reg-reg	ROUNDPx, ROUNDSx, VCVTPS2PH, VRNDSCALExx	All SIMD operands, vector lengths

The static rounding-mode override in AVX-512 also implies the “suppress-all-exceptions” (SAE) attribute. The SAE effect is as if all the MXCSR mask bits are set, and none of the MXCSR flags will be updated. Using static rounding-mode via EVEX without SAE is not supported.

Static Rounding Mode and SAE control can be enabled in the encoding of the instruction by setting the EVEX.b bit to 1 in a register-register vector instruction. In such a case, vector length is assumed to be MAX_VL (512-bit in case of AVX-512 packed vector instructions) or 128-bit for scalar instructions. Table 2-5 summarizes the possible static rounding-mode assignments in AVX-512 instructions.

Note that some instructions already allow to specify the rounding mode statically via immediate bits. In such case, the immediate bits take precedence over the embedded rounding mode (in the same vein that they take precedence over whatever MXCSR.RM says).

Table 2-5. Static Rounding Mode

Function	Description
{rn-sae}	Round to nearest (even) + SAE
{rd-sae}	Round down (toward -inf) + SAE
{ru-sae}	Round up (toward +inf) + SAE
{rz-sae}	Round toward zero (Truncate) + SAE

An example of use would be in the following instructions:

```
vaddps zmm7 {k6}, zmm2, zmm4, {rd-sae}
```

Which would perform the single-precision floating-point addition of vectors zmm2 and zmm4 with round-towards-minus-infinity, leaving the result in vector zmm7 using k6 as conditional writemask.

Note that MXCSR.RM bits are ignored and unaffected by the outcome of this instruction.

Examples of instructions instances where the static rounding-mode is not allowed would be:

```
; rounding-mode already specified in the instruction immediate
```

```
vrndscaleps zmm7 {k6}, zmm2, 0x00
```

```
; instructions with memory operands
```

```
vmulps zmm7 {k6}, zmm2, [rax], {rd-sae}
```

2.5.5 Compressed Disp8*N Encoding

EVEX encoding supports a new displacement representation that allows for a more compact encoding of memory addressing commonly used in unrolled code, where an 8-bit displacement can address a range exceeding the dynamic range of an 8-bit value. This compressed displacement encoding is referred to as disp8*N, where N is a constant implied by the memory operation characteristic of each instruction.

The compressed displacement is based on the assumption that the effective displacement (of a memory operand occurring in a loop) is a multiple of the granularity of the memory access of each iteration. Since the Base register in memory addressing already provides byte-granular resolution, the lower bits of the traditional disp8 operand becomes redundant, and can be implied from the memory operation characteristic.

The memory operation characteristics depend on the following:

- The destination operand is updated as a full vector, a single element, or multi-element tuples.
- The memory source operand (or vector source operand if the destination operand is memory) is fetched (or treated) as a full vector, a single element, or multi-element tuples.

For example,

```
vaddps zmm7, zmm2, disp8[membase + index*8]
```

The destination zmm7 is updated as a full 512-bit vector, and 64-bytes of data are fetched from memory as a full vector; the next unrolled iteration may fetch from memory in 64-byte granularity per iteration. There are 6 bits of lowest address that can be compressed, hence $N = 2^6 = 64$. The contribution of “disp8” to effective address calculation is $64 * \text{disp8}$.


```
vbroadcastf32x4 zmm7, disp8[membase + index*8]
```

In VBROADCASTF32x4, memory is fetched as a 4tuple of 4 32-bit entities. Hence the common lowest address bits that can be compressed is 4, corresponding to the 4tuple width of $2^4 = 16$ bytes (4x32 bits). Therefore, $N = 2^4$.

For EVEX encoded instructions that update only one element in the destination, or source element is fetched individually, the number of lowest address bits that can be compressed is generally the width in bytes of the data element, hence $N = 2^{(\text{width})}$.

2.6 MEMORY ALIGNMENT

Memory alignment requirements on EVEX-encoded SIMD instructions are similar to VEX-encoded SIMD instructions. Memory alignment applies to EVEX-encoded SIMD instructions in three categories:

- Explicitly-aligned SIMD load and store instructions accessing 64 bytes of memory with EVEX prefix encoded vector length of 512 bits (e.g., VMOVAPD, VMOVAPS, VMOVDQA, etc.). These instructions always require memory address to be aligned on 64-byte boundary.
- Explicitly-unaligned SIMD load and store instructions accessing 64 bytes or less of data from memory (e.g. VMOVUPD, VMOVUPS, VMOVDQU, VMOVQ, VMOVD, etc.). These instructions do not require memory address to be aligned on natural vector-length byte boundary.
- Most arithmetic and data processing instructions encoded using EVEX support memory access semantics. When these instructions access from memory, there are no alignment restrictions.

Software may see performance penalties when unaligned accesses cross cacheline boundaries or vector-length naturally-aligned boundaries, so reasonable attempts to align commonly used data sets should continue to be pursued.

Atomic memory operation in Intel 64 and IA-32 architecture is guaranteed only for a subset of memory operand sizes and alignment scenarios. The guaranteed atomic operations are described in Section 7.1.1 of IA-32 Intel® Architecture Software Developer's Manual, Volumes 3A. AVX and FMA instructions do not introduce any new guaranteed atomic memory operations.

AVX-512 instructions may generate an #AC(0) fault on misaligned 4 or 8-byte memory references in Ring-3 when CRO.AM=1. 16, 32 and 64-byte memory references will not generate #AC(0) fault. See Table 2-7 for details.

Certain AVX-512 Foundation instructions always require 64-byte alignment (see the complete list of VEX and EVEX encoded instructions in Table 2-6). These instructions will #GP(0) if not aligned to 64-byte boundaries.

Table 2-6. SIMD Instructions Requiring Explicitly Aligned Memory

Require 16-byte alignment	Require 32-byte alignment	Require 64-byte alignment*
(V)MOVDQA xmm, m128	VMOVDQA ymm, m256	VMOVDQA zmm, m512
(V)MOVDQA m128, xmm	VMOVDQA m256, ymm	VMOVDQA m512, zmm
(V)MOVAPS xmm, m128	VMOVAPS ymm, m256	VMOVAPS zmm, m512
(V)MOVAPS m128, xmm	VMOVAPS m256, ymm	VMOVAPS m512, zmm
(V)MOVAPD xmm, m128	VMOVAPD ymm, m256	VMOVAPD zmm, m512
(V)MOVAPD m128, xmm	VMOVAPD m256, ymm	VMOVAPD m512, zmm
(V)MOVNTDQA xmm, m128	VMOVNTPS m256, ymm	VMOVNTPS m512, zmm
(V)MOVNTPS m128, xmm	VMOVNTPD m256, ymm	VMOVNTPD m512, zmm
(V)MOVNTPD m128, xmm	VMOVNTDQ m256, ymm	VMOVNTDQ m512, zmm
(V)MOVNTDQ m128, xmm	VMOVNTDQA ymm, m256	VMOVNTDQA zmm, m512

Table 2-7. Instructions Not Requiring Explicit Memory Alignment

(V)MOVDQU xmm, m128	VMOVDQU ymm, m256	VMOVDQU zmm, m512
(V)MOVDQU m128, m128	VMOVDQU m256, ymm	VMOVDQU m512, zmm
(V)MOVUPS xmm, m128	VMOVUPS ymm, m256	VMOVUPS zmm, m512
(V)MOVUPS m128, xmm	VMOVUPS m256, ymm	VMOVUPS m512, zmm
(V)MOVUPD xmm, m128	VMOVUPD ymm, m256	VMOVUPD zmm, m512
(V)MOVUPD m128, xmm	VMOVUPD m256, ymm	VMOVUPD m512, zmm

2.7 SIMD FLOATING-POINT EXCEPTIONS

AVX-512 instructions can generate SIMD floating-point exceptions (#XM) if embedded “suppress all exceptions” (SAE) in EVEX is not set. When SAE is not set, these instructions will respond to exception masks of MXCSR in the same way as VEX-encoded AVX instructions. When CR4.OSXMMEXCPT=0 any unmasked FP exceptions generate an Undefined Opcode exception (#UD).

2.8 INSTRUCTION EXCEPTION SPECIFICATION

Exception behavior of VEX-encoded Intel AVX and Intel AVX2 instructions are described in *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2A*. Exception behavior of AVX-512 Foundation instructions and additional 512-bit extensions are described in Section 4.10, “Exception Classifications of EVEX-Encoded instructions” and Section 4.11, “Exception Classifications of Opmask instructions”.

2.9 CPUID INSTRUCTION

CPUID—CPU Identification

Opcode	Instruction	64-Bit Mode	Compat/ Leg Mode	Description
OF A2	CPUID	Valid	Valid	Returns processor identification and feature information to the EAX, EBX, ECX, and EDX registers, as determined by input entered in EAX (in some cases, ECX as well).

Description

The ID flag (bit 21) in the EFLAGS register indicates support for the CPUID instruction. If a software procedure can set and clear this flag, the processor executing the procedure supports the CPUID instruction. This instruction operates the same in non-64-bit modes and 64-bit mode.

CPUID returns processor identification and feature information in the EAX, EBX, ECX, and EDX registers.¹ The instruction's output is dependent on the contents of the EAX register upon execution (in some cases, ECX as well). For example, the following pseudocode loads EAX with 00H and causes CPUID to return a Maximum Return Value and the Vendor Identification String in the appropriate registers:

```
MOV EAX, 00H
CPUID
```

Table 2-8 shows information returned, depending on the initial value loaded into the EAX register. Table 2-9 shows the maximum CPUID input value recognized for each family of IA-32 processors on which CPUID is implemented.

Two types of information are returned: basic and extended function information. If a value is entered for CPUID.EAX is invalid for a particular processor, the data for the highest basic information leaf is returned. For example, using the Intel Core 2 Duo E6850 processor, the following is true:

```
CPUID.EAX = 05H (* Returns MONITOR/MWAIT leaf. *)
CPUID.EAX = 0AH (* Returns Architectural Performance Monitoring leaf. *)
CPUID.EAX = 0BH (* INVALID: Returns the same information as CPUID.EAX = 0AH. *)
CPUID.EAX = 80000008H (* Returns virtual/physical address size data. *)
CPUID.EAX = 8000000AH (* INVALID: Returns same information as CPUID.EAX = 0AH. *)
```

When CPUID returns the highest basic leaf information as a result of an invalid input EAX value, any dependence on input ECX value in the basic leaf is honored.

CPUID can be executed at any privilege level to serialize instruction execution. Serializing instruction execution guarantees that any modifications to flags, registers, and memory for previous instructions are completed before the next instruction is fetched and executed.

See also:

"Serializing Instructions" in Chapter 8, "Multiple-Processor Management," in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*

"Caching Translation Information" in Chapter 4, "Paging," in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*.

1. On Intel 64 processors, CPUID clears the high 32 bits of the RAX/RBX/RCX/RDX registers in all modes.

Table 2-8. Information Returned by CPUID Instruction

Initial EAX Value	Information Provided about the Processor	
Basic CPUID Information		
0H	EAX EBX ECX EDX	Maximum Input Value for Basic CPUID Information (see Table 2-9) "Genu" "ntel" "inel"
01H	EAX EBX ECX EDX	Version Information: Type, Family, Model, and Stepping ID (see Figure 2-4) Bits 7-0: Brand Index Bits 15-8: CLFLUSH line size (Value * 8 = cache line size in bytes) Bits 23-16: Maximum number of addressable IDs for logical processors in this physical package*. Bits 31-24: Initial APIC ID Feature Information (see Figure 2-5 and Table 2-11) Feature Information (see Figure 2-6 and Table 2-12) NOTES: * The nearest power-of-2 integer that is not smaller than EBX[23:16] is the maximum number of unique initial APIC IDs reserved for addressing different logical processors in a physical package.
02H	EAX EBX ECX EDX	Cache and TLB Information (see Table 2-13) Cache and TLB Information Cache and TLB Information Cache and TLB Information
03H	EAX EBX ECX EDX	Reserved. Reserved. Bits 00-31 of 96 bit processor serial number. (Available in Pentium III processor only; otherwise, the value in this register is reserved.) Bits 32-63 of 96 bit processor serial number. (Available in Pentium III processor only; otherwise, the value in this register is reserved.) NOTES: Processor serial number (PSN) is not supported in the Pentium 4 processor or later. On all models, use the PSN flag (returned using CPUID) to check for PSN support before accessing the feature.
CPUID leaves > 3 < 80000000 are visible only when IA32_MISC_ENABLES.BOOT_NT4[bit 22] = 0 (default).		
Deterministic Cache Parameters Leaf		
04H	EAX	NOTES: Leaf 04H output depends on the initial value in ECX. See also: "INPUT EAX = 4: Returns Deterministic Cache Parameters for each level on page 2-35." Bits 4-0: Cache Type Field 0 = Null - No more caches 1 = Data Cache 2 = Instruction Cache 3 = Unified Cache 4-31 = Reserved

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
		<p>Bits 7-5: Cache Level (starts at 1) Bits 8: Self Initializing cache level (does not need SW initialization) Bits 9: Fully Associative cache</p> <p>Bits 13-10: Reserved Bits 25-14: Maximum number of addressable IDs for logical processors sharing this cache*, ** Bits 31-26: Maximum number of addressable IDs for processor cores in the physical package*, ***, ****</p> <p>EBX Bits 11-00: L = System Coherency Line Size* Bits 21-12: P = Physical Line partitions* Bits 31-22: W = Ways of associativity*</p> <p>ECX Bits 31-00: S = Number of Sets*</p> <p>EDX Bit 0: WBINVD/INVD behavior on lower level caches Bit 10: Write-Back Invalidate/Invalidate 0 = WBINVD/INVD from threads sharing this cache acts upon lower level caches for threads sharing this cache 1 = WBINVD/INVD is not guaranteed to act upon lower level caches of non-originating threads sharing this cache. Bit 1: Cache Inclusiveness 0 = Cache is not inclusive of lower cache levels. 1 = Cache is inclusive of lower cache levels. Bit 2: Complex cache indexing 0 = Direct mapped cache 1 = A complex function is used to index the cache, potentially using all address bits. Bits 31-03: Reserved = 0</p> <p>NOTES: * Add one to the return value to get the result. ** The nearest power-of-2 integer that is not smaller than (1 + EAX[25:14]) is the number of unique initial APIC IDs reserved for addressing different logical processors sharing this cache *** The nearest power-of-2 integer that is not smaller than (1 + EAX[31:26]) is the number of unique Core_IDs reserved for addressing different processor cores in a physical package. Core ID is a subset of bits of the initial APIC ID. ****The returned value is constant for valid initial values in ECX. Valid ECX values start from 0.</p>
	<i>MONITOR/MWAIT Leaf</i>	
05H	EAX EBX ECX	<p>Bits 15-00: Smallest monitor-line size in bytes (default is processor's monitor granularity) Bits 31-16: Reserved = 0</p> <p>Bits 15-00: Largest monitor-line size in bytes (default is processor's monitor granularity) Bits 31-16: Reserved = 0</p> <p>Bits 00: Enumeration of Monitor-Mwait extensions (beyond EAX and EBX registers) supported Bits 01: Supports treating interrupts as break-event for MWAIT, even when interrupts disabled Bits 31 - 02: Reserved</p>

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
	EDX	<p>Bits 03 - 00: Number of C0* sub C-states supported using MWait Bits 07 - 04: Number of C1* sub C-states supported using MWait Bits 11 - 08: Number of C2* sub C-states supported using MWait Bits 15 - 12: Number of C3* sub C-states supported using MWait Bits 19 - 16: Number of C4* sub C-states supported using MWait Bits 23 - 20: Number of C5* sub C-states supported using MWait Bits 27 - 24: Number of C6* sub C-states supported using MWait Bits 31 - 28: Number of C7* sub C-states supported using MWait</p> <p>NOTE: * The definition of C0 through C7 states for MWait extension are processor-specific C-states, not ACPI C-states.</p>
<i>Thermal and Power Management Leaf</i>		
06H	EAX	<p>Bit 00: Digital temperature sensor is supported if set Bit 01: Intel Turbo Boost Technology Available (see description of IA32_MISC_ENABLE[38]). Bit 02: ARAT. APIC-Timer-always-running feature is supported if set.</p>
	EBX	<p>Bit 03: Reserved Bit 04: PLN. Power limit notification controls are supported if set. Bit 05: ECMD. Clock modulation duty cycle extension is supported if set. Bit 06: PTM. Package thermal management is supported if set. Bit 07: HWP. HWP base registers (IA32_PM_ENABLE[bit 0], IA32_HWP_CAPABILITIES, IA32_HWP_REQUEST, IA32_HWP_STATUS) are supported if set. Bit 08: HWP_Notification. IA32_HWP_INTERRUPT MSR is supported if set. Bit 09: HWP_Activity_Window. IA32_HWP_REQUEST[bits 41:32] is supported if set. Bit 10: HWP_Energy_Performance_Preference. IA32_HWP_REQUEST[bits 31:24] is supported if set. Bit 11: HWP_Package_Level_Request. IA32_HWP_REQUEST_PKG MSR is supported if set. Bit 12: Reserved. Bit 13: HDC. HDC base registers IA32PKG_HDC_CTL, IA32PM_CTL1, IA32_THREAD_STALL MSRs are supported if set. Bits 31 - 15: Reserved Bits 03 - 00: Number of Interrupt Thresholds in Digital Thermal Sensor Bits 31 - 04: Reserved</p>
	ECX	<p>Bit 00: Hardware Coordination Feedback Capability (Presence of IA32_MPERF and IA32_APERF). The capability to provide a measure of delivered processor performance (since last reset of the counters), as a percentage of the expected processor performance when running at the TSC frequency. Bits 02 - 01: Reserved = 0 Bit 03: The processor supports performance-energy bias preference if CPUID.06H:ECX.SETBH[bit 3] is set and it also implies the presence of a new architectural MSR called IA32_ENERGY_PERF_BIAS (1B0H) Bits 31 - 04: Reserved = 0</p>
	EDX	Reserved = 0
<i>Structured Extended Feature Flags Enumeration Leaf (Output depends on ECX input value)</i>		
07H		<p>NOTES: Leaf 07H main leaf (ECX = 0). If ECX contains an invalid sub-leaf index, EAX/EBX/ECX/EDX return 0.</p>
	EAX	Bits 31-00: Reports the maximum number sub-leaves that are supported in leaf 07H.

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
	EBX	Bit 00: FSGSBASE. Supports RDFSBASE/RDGSBASE/WRFSBASE/WRGSBASE if 1. Bit 01: IA32_TSC_ADJUST MSR is supported if 1. Bit 02: SGX Bit 03: BMI1 Bit 04: HLE Bit 05: AVX2 Bit 07: SMEP. Supports Supervisor Mode Execution Protection if 1. Bit 06: Reserved Bit 08: BMI2 Bit 09: ERMS Bit 10: INVPCID Bit 11: RTM Bit 12: Supports Platform Quality of Service Monitoring (PQM) capability if 1. Bit 13: Deprecates FPU CS and FPU DS values if 1. Bit 14: Intel Memory Protection Extensions Bit 15: Supports Platform Quality of Service Enforcement (PQE) capability if 1. Bit 16: AVX512F Bit 17: AVX512DQ Bit 18: RDSEED Bit 19: ADX Bit 20: SMAP Bit 21: AVX512IFMA Bit 22: Reserved Bit 23: CLFLUSHOPT Bit 24: CLWB Bit 25: Intel Processor Trace Bit 26: AVX512PF Bit 27: AVX512ER Bit 28: AVX512CD Bit 29: SHA Bit 30: AVX512BW Bit 31: AVX512VL
	ECX	Bit 00: PREFETCHWT1 Bit 01: AVX512VBMI Bit 02: UMIP. Supports user-mode instruction prevention if 1. Bit 03: PKU. Supports protection keys for user-mode pages if 1. Bit 04: OSPKE. If 1, OS has set CR4.PKE to enable protection keys (and the RDPKRU/WRPKRU instructions). Bits 13:05: Reserved Bit 14: AVX512_VPOPCNTDQ Bits 16:15: Reserved Bits 21 - 17: The value of MAWAU used by the BNDLDX and BNDSTX instructions in 64-bit mode. Bit 22: RDPID. Supports Read Processor ID if 1. Bits 29 - 23: Reserved. Bit 30: SGX_LC. Supports SGX Launch Configuration if 1. Bit 31: Reserved.
	EDX	Bits 01 - 00: Reserved Bit 02: AVX512_4VNNIW (Vector instructions for deep learning enhanced word variable precision.) Bit 03: AVX512_4FMAPS (Vector instructions for deep learning floating-point single precision.) Bits 31-04: Reserved

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
Structured Extended Feature Enumeration Sub-leaves (EAX = 07H, ECX = n, n ≥ 1)		
07H	NOTES: Leaf 07H output depends on the initial value in ECX. If ECX contains an invalid sub leaf index, EAX/EBX/ECX/EDX return 0. EAX This field reports 0 if the sub-leaf index, n, is invalid*; otherwise it is reserved. EBX This field reports 0 if the sub-leaf index, n, is invalid*; otherwise it is reserved. ECX This field reports 0 if the sub-leaf index, n, is invalid*; otherwise it is reserved. EDX This field reports 0 if the sub-leaf index, n, is invalid*; otherwise it is reserved.	
Direct Cache Access Information Leaf		
09H	EAX EBX ECX EDX	Value of bits [31:0] of IA32_PLATFORM_DCA_CAP MSR (address 1F8H) Reserved Reserved Reserved
Architectural Performance Monitoring Leaf		
0AH	EAX EBX ECX EDX	Bits 07 - 00: Version ID of architectural performance monitoring Bits 15- 08: Number of general-purpose performance monitoring counter per logical processor Bits 23 - 16: Bit width of general-purpose, performance monitoring counter Bits 31 - 24: Length of EBX bit vector to enumerate architectural performance monitoring events Bit 00: Core cycle event not available if 1 Bit 01: Instruction retired event not available if 1 Bit 02: Reference cycles event not available if 1 Bit 03: Last-level cache reference event not available if 1 Bit 04: Last-level cache misses event not available if 1 Bit 05: Branch instruction retired event not available if 1 Bit 06: Branch mispredict retired event not available if 1 Bits 31- 07: Reserved = 0 Reserved = 0 Bits 04 - 00: Number of fixed-function performance counters (if Version ID > 1) Bits 12- 05: Bit width of fixed-function performance counters (if Version ID > 1) Reserved = 0
Extended Topology Enumeration Leaf		
0BH	NOTES: Most of Leaf 0BH output depends on the initial value in ECX. The EDX output of leaf 0BH is always valid and does not vary with input value in ECX. Output value in ECX[7:0] always equals input value in ECX[7:0]. For sub-leaves that returns an invalid level-type of 0 in ECX[15:8]; EAX and EBX will return 0. If an input value N in ECX returns the invalid level-type of 0 in ECX[15:8], other input values with ECX > N also return 0 in ECX[15:8] EAX Bits 04-00: Number of bits to shift right on x2APIC ID to get a unique topology ID of the next level type*. All logical processors with the same next level ID share current level. Bits 31-5: Reserved. EBX Bits 15 - 00: Number of logical processors at this level type. The number reflects configuration as shipped by Intel**. Bits 31- 16: Reserved.	

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
	ECX	Bits 07 - 00: Level number. Same value in ECX input Bits 15 - 08: Level type***. Bits 31 - 16: Reserved.
	EDX	Bits 31 - 00: x2APIC ID the current logical processor. NOTES: * Software should use this field (EAX[4:0]) to enumerate processor topology of the system. ** Software must not use EBX[15:0] to enumerate processor topology of the system. This value in this field (EBX[15:0]) is only intended for display/diagnostic purposes. The actual number of logical processors available to BIOS/OS/Applications may be different from the value of EBX[15:0], depending on software and platform hardware configurations. *** The value of the "level type" field is not related to level numbers in any way, higher "level type" values do not mean higher levels. Level type field has the following encoding: 0: invalid 1: SMT 2: Core 3-255: Reserved
<i>Processor Extended State Enumeration Main Leaf (EAX = 0DH, ECX = 0)</i>		
0DH		NOTES: Leaf 0DH main leaf (ECX = 0).
	EAX	Bits 31-00: Reports the valid bit fields of the lower 32 bits of the XFEATURE_ENABLED_MASK register. If a bit is 0, the corresponding bit field in XCRO is reserved. Bit 00: legacy x87 Bit 01: 128-bit SSE Bit 02: 256-bit AVX Bits 04 - 03: MPX state Bit 07 - 05: AVX-512 state Bit 08: Used for IA32_XSS Bit 09: PKRU state Bits 31-10: Reserved
	EBX	Bits 31-00: Maximum size (bytes, from the beginning of the XSAVE/XRSTOR save area) required by enabled features in XCRO. May be different than ECX if some features at the end of the XSAVE save area are not enabled.
	ECX	Bit 31-00: Maximum size (bytes, from the beginning of the XSAVE/XRSTOR save area) of the XSAVE/XRSTOR save area required by all supported features in the processor, i.e all the valid bit fields in XCRO.
	EDX	Bit 31-0: Reports the valid bit fields of the upper 32 bits of the XCRO register. If a bit is 0, the corresponding bit field in XCRO is reserved

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
Processor Extended State Enumeration Sub-leaf (EAX = 0DH, ECX = 1)		
0DH	EAX	Bit 00: XSAVEOPT is available Bit 01: Supports XSAVEC and the compacted form of XRSTOR if set Bit 02: Supports XGETBV with ECX = 1 if set Bit 03: Supports XSAVES/XRSTORS and IA32_XSS if set Bits 31-04: Reserved
	EBX	Bits 31-00: The size in bytes of the XSAVE area containing all states enabled by XCRO IA32_XSS.
	ECX	Bits 31-00: Reports the supported bits of the lower 32 bits of the IA32_XSS MSR. IA32_XSS[n] can be set to 1 only if ECX[n] is 1. Bits 07-00: Used for XCRO Bit 08: PT state Bit 09: Used for XCRO Bits 31-10: Reserved
	EDX	Bits 31-00: Reports the supported bits of the upper 32 bits of the IA32_XSS MSR. IA32_XSS[n+32] can be set to 1 only if EDX[n] is 1. Bits 31-00: Reserved
Processor Extended State Enumeration Sub-leaves (EAX = 0DH, ECX = n, n > 1)		
0DH		NOTES: Leaf 0DH output depends on the initial value in ECX. Each sub-leaf index (starting at position 2) is supported if it corresponds to a supported bit in either the XCRO register or the IA32_XSS MSR. * If ECX contains an invalid sub-leaf index, EAX/EBX/ECX/EDX return 0. Sub-leaf n (0 ≤ n ≤ 31) is invalid if sub-leaf 0 returns 0 in EAX[n] and sub-leaf 1 returns 0 in ECX[n]. Sub-leaf n (32 ≤ n ≤ 63) is invalid if sub-leaf 0 returns 0 in EDX[n-32] and sub-leaf 1 returns 0 in EDX[n-32].
	EAX	Bits 31-0: The size in bytes (from the offset specified in EBX) of the save area for an extended state feature associated with a valid sub-leaf index, n. This field reports 0 if the sub-leaf index, n, is invalid*.
	EBX	Bits 31-0: The offset in bytes of this extended state component’s save area from the beginning of the XSAVE/XRSTOR area. This field reports 0 if the sub-leaf index, n, does not map to a valid bit in the XCRO register*.
	ECX	Bit 0 is set if the bit n (corresponding to the sub-leaf index) is supported in the IA32_XSS MSR; it is clear if bit n is instead supported in XCRO. Bit 1 is set if, when the compacted format of an XSAVE area is used, this extended state component located on the next 64-byte boundary following the preceding state component (otherwise, it is located immediately following the preceding state component). Bits 31:02 are reserved. This field reports 0 if the sub-leaf index, n, is invalid*.
	EDX	This field reports 0 if the sub-leaf index, n, is invalid*; otherwise it is reserved.
Platform QoS Monitoring Enumeration Sub-leaf (EAX = 0FH, ECX = 0)		
0FH		NOTES: Leaf 0FH output depends on the initial value in ECX. Sub-leaf index 0 reports valid resource type starting at bit position 1 of EDX.
	EAX	Reserved.
	EBX	Bits 31-0: Maximum range (zero-based) of RMID within this physical processor of all types.
	ECX	Reserved.
	EDX	Bit 00: Reserved. Bit 01: Supports L3 Cache QoS Monitoring if 1. Bits 31 - 02: Reserved

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
L3 Cache QoS Monitoring Capability Enumeration Sub-leaf (EAX = 0FH, ECX = 1)		
0FH	NOTES: Leaf 0FH output depends on the initial value in ECX.	
EAX	Reserved.	
EBX	Bits 31-0: Conversion factor from reported IA32_QM_CTR value to occupancy metric (bytes).	
ECX	Maximum range (zero-based) of RMID of this resource type.	
EDX	Bit 00: Supports L3 occupancy monitoring if 1. Bits 31:01: Reserved	
Platform QoS Enforcement Enumeration Sub-leaf (EAX = 10H, ECX = 0)		
10H	NOTES: Leaf 10H output depends on the initial value in ECX. Sub-leaf index 0 reports valid resource identification (ResID) starting at bit position 1 of EBX.	
EAX	Reserved.	
EBX	Bit 00: Reserved. Bit 01: Supports L3 Cache QoS Enforcement if 1. Bits 31 - 02: Reserved.	
ECX	Reserved.	
EDX	Reserved.	
L3 Cache QoS Enforcement Enumeration Sub-leaf (EAX = 10H, ECX = ResID =1)		
10H	NOTES: Leaf 10H output depends on the initial value in ECX.	
EAX	Bits 4:0: Length of the capacity bit mask for the corresponding ResID. Bits 31:05: Reserved	
EBX	Bits 31-0: Bit-granular map of isolation/contention of allocation units.	
ECX	Bit 00: Reserved. Bit 01: Updates of COS should be infrequent if 1. Bit 02: Code and Data Prioritization Technology supported if 1. Bits 31:03: Reserved	
EDX	Bits 15:0: Highest COS number supported for this ResID. Bits 31:16: Reserved	
Intel Processor Trace Enumeration Main Leaf (EAX = 14H, ECX = 0)		
14H	NOTES: Leaf 14H main leaf (ECX = 0).	
EAX	Bits 31-0: Reports the maximum sub-leaf supported in leaf 14H.	
EBX	Bit 00: If 1, Indicates that IA32_RTIT_CTL.CR3Filter can be set to 1, and that IA32_RTIT_CR3_MATCH MSR can be accessed. Bits 01: If 1, Indicates support of Configurable PSB and Cycle-Accurate Mode. Bits 02: If 1, Indicates support of IP Filtering, TraceStop filtering, and preservation of Intel PT MSRs across warm reset. Bits 03: If 1, Indicates support of MTC timing packet and suppression of COFI-based packets. Bits 31: 04: Reserved	

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
	ECX	Bit 00: If 1, Tracing can be enabled with IA32_RTIT_CTL.ToPA = 1, hence utilizing the ToPA output scheme; IA32_RTIT_OUTPUT_BASE and IA32_RTIT_OUTPUT_MASK_PTRS MSRs can be accessed. Bit 01: If 1, ToPA tables can hold any number of output entries, up to the maximum allowed by the MaskOrTableOffset field of IA32_RTIT_OUTPUT_MASK_PTRS. Bits 02: If 1, Indicates support of Single-Range Output scheme. Bits 03: If 1, Indicates support of output to Trace Transport subsystem. Bit 30:04: Reserved Bit 31: If 1, Generated packets which contain IP payloads have LIP values, which include the CS base component.
	EDX	Bits 31- 00: Reserved
<i>Intel Processor Trace Enumeration Sub-leaf (EAX = 14H, ECX = 1)</i>		
14H	EAX	Bits 2:0: Number of configurable Address Ranges for filtering. Bits 15-03: Reserved Bit 31:16: Bitmap of supported MTC period encodings
	EBX	Bits 15-0: Bitmap of supported Cycle Threshold value encodings Bit 31:16: Bitmap of supported Configurable PSB frequency encodings
	ECX	Bits 31-00: Reserved
	EDX	Bits 31- 00: Reserved
<i>Time Stamp Counter and Core Crystal Clock Information Leaf</i>		
15H	NOTES: If EBX[31:0] is 0, the TSC and “core crystal clock” ratio is not enumerated. EBX[31:0]/EAX[31:0] indicates the ratio of the TSC frequency and the core crystal clock frequency. If ECX is 0, the core crystal clock frequency is not enumerated. “TSC frequency” = “core crystal clock frequency” * EBX/EAX. The core crystal clock may differ from the reference clock, bus clock, or core clock frequencies.	
	EAX	Bits 31:0: An unsigned integer which is the denominator of the TSC/“core crystal clock” ratio.
	EBX	Bits 31:0: An unsigned integer which is the numerator of the TSC/“core crystal clock” ratio.
	ECX	Bits 31:0: An unsigned integer which is the nominal frequency of the core crystal clock in Hz.
	EDX	Bits 31:0: Reserved = 0.

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
Processor Frequency Information Leaf		
16H	EAX EBX ECX EDX	Bits 15:0: Processor Base Frequency (in MHz). Bits 31:16: Reserved =0 Bits 15:0: Maximum Frequency (in MHz). Bits 31:16: Reserved = 0 Bits 15:0: Bus (Reference) Frequency (in MHz). Bits 31:16: Reserved = 0 Reserved NOTES: * Data is returned from this interface in accordance with the processor’s specification and does not reflect actual values. Suitable use of this data includes the display of processor information in like manner to the processor brand string and for determining the appropriate range to use when displaying processor information e.g. frequency history graphs. The returned information should not be used for any other purpose as the returned information does not accurately correlate to information / counters returned by other processor interfaces. While a processor may support the Processor Frequency Information leaf, fields that return a value of zero are not supported.
System-On-Chip Vendor Attribute Enumeration Main Leaf (EAX = 17H, ECX = 0)		
17H		NOTES: Leaf 17H main leaf (ECX = 0). Leaf 17H output depends on the initial value in ECX. Leaf 17H sub-leaves 1 through 3 reports SOC Vendor Brand String. Leaf 17H is valid if MaxSOCID_Index >= 3. Leaf 17H sub-leaves 4 and above are reserved.
	EAX EBX ECX EDX	Bits 31 - 00: MaxSOCID_Index. Reports the maximum input value of supported sub-leaf in leaf 17H. Bits 15 - 00: SOC Vendor ID. Bit 16: IsVendorScheme. If 1, the SOC Vendor ID field is assigned via an industry standard enumeration scheme. Otherwise, the SOC Vendor ID field is assigned by Intel. Bits 31 - 17: Reserved = 0. Bits 31 - 00: Project ID. A unique number an SOC vendor assigns to its SOC projects. Bits 31 - 00: Stepping ID. A unique number within an SOC project that an SOC vendor assigns.
System-On-Chip Vendor Attribute Enumeration Sub-leaf (EAX = 17H, ECX = 1..3)		
17H	EAX EBX ECX EDX	Bit 31 - 00: SOC Vendor Brand String. UTF-8 encoded string. Bit 31 - 00: SOC Vendor Brand String. UTF-8 encoded string. Bit 31 - 00: SOC Vendor Brand String. UTF-8 encoded string. Bit 31 - 00: SOC Vendor Brand String. UTF-8 encoded string. NOTES: Leaf 17H output depends on the initial value in ECX. SOC Vendor Brand String is a UTF-8 encoded string padded with trailing bytes of 00H. The complete SOC Vendor Brand String is constructed by concatenating in ascending order of EAX:EBX:ECX:EDX and from the sub-leaf 1 fragment towards sub-leaf 3.

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor
<i>System-On-Chip Vendor Attribute Enumeration Sub-leaves (EAX = 17H, ECX > MaxSOCID_Index)</i>	
17H	<p>NOTES: Leaf 17H output depends on the initial value in ECX.</p> <p>EAX Bits 31 - 00: Reserved = 0. EBX Bits 31 - 00: Reserved = 0. ECX Bits 31 - 00: Reserved = 0. EDX Bits 31 - 00: Reserved = 0.</p>
<i>Deterministic Address Translation Parameters Main Leaf (EAX = 18H, ECX = 0)</i>	
18H	<p>NOTES: Each sub-leaf enumerates a different address translations structure. Valid sub-leaves do not need to be contiguous or in any particular order. A valid sub-leaf may be in a higher input ECX value than an invalid sub-leaf or than a valid sub-leaf of a higher or lower-level structure. If ECX contains an invalid sub-leaf index, EAX/EBX/ECX/EDX return 0. Sub-leaf index n is invalid if n exceeds the value that sub-leaf 0 returns in EAX. * Add one to the return value to get the result.</p> <p>EAX Bits 31 - 00: Reports the maximum input value of supported sub-leaf in leaf 18H. EBX Bit 00: 4K page size entries supported by this structure. Bit 01: 2MB page size entries supported by this structure. Bit 02: 4MB page size entries supported by this structure. Bit 03: 1 GB page size entries supported by this structure. Bits 07 - 04: Reserved. Bits 10 - 08: Partitioning (0: Soft partitioning between the logical processors sharing this structure). Bits 15 - 11: Reserved. Bits 31 - 16: W = Ways of associativity.</p> <p>ECX Bits 31 - 00: S = Number of Sets.</p> <p>EDX Bits 04 - 00: Translation cache type field. 0000b: Null (indicates this sub-leaf is not valid). 0001b: Data TLB. 0010b: Instruction TLB. 0011b: Unified TLB. All other encodings are reserved. Bits 07 - 05: Translation cache level (starts at 1). Bit 08: Fully associative structure. Bits 13 - 09: Reserved. Bits 25- 14: Maximum number of addressable IDs for logical processors sharing this translation cache* Bits 31 - 26: Reserved.</p>
<i>Deterministic Address Translation Parameters Sub-leaf (EAX = 18H, ECX ≥ 1)</i>	
18H	<p>NOTES: If ECX contains an invalid sub-leaf index, EAX/EBX/ECX/EDX return 0. Sub-leaf index n is invalid if n exceeds the value that sub-leaf 0 returns in EAX. * Add one to the return value to get the result.</p> <p>EAX Bits 31 - 00: Reserved.</p>

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
	EBX	Bit 00: 4K page size entries supported by this structure. Bit 01: 2MB page size entries supported by this structure. Bit 02: 4MB page size entries supported by this structure. Bit 03: 1 GB page size entries supported by this structure. Bits 07 - 04: Reserved. Bits 10 - 08: Partitioning (0: Soft partitioning between the logical processors sharing this structure). Bits 15 - 11: Reserved. Bits 31 - 16: W = Ways of associativity.
	ECX	Bits 31 - 00: S = Number of Sets.
	EDX	Bits 04 - 00: Translation cache type field. 0000b: Null (indicates this sub-leaf is not valid). 0001b: Data TLB. 0010b: Instruction TLB. 0011b: Unified TLB. All other encodings are reserved. Bits 07 - 05: Translation cache level (starts at 1). Bit 08: Fully associative structure. Bits 13 - 09: Reserved. Bits 25- 14: Maximum number of addressable IDs for logical processors sharing this translation cache* Bits 31 - 26: Reserved.
<i>Unimplemented CPUID Leaf Functions</i>		
40000000H - 4FFFFFFFH		Invalid. No existing or future CPU will return processor identification or feature information if the initial EAX value is in the range 40000000H to 4FFFFFFFH.
<i>Extended Function CPUID Information</i>		
80000000H	EAX	Maximum Input Value for Extended Function CPUID Information (see Table 2-9).
	EBX	Reserved
	ECX	Reserved
	EDX	Reserved
80000001H	EAX	Extended Processor Signature and Feature Bits.
	EBX	Reserved
	ECX	Bit 0: LAHF/SAHF available in 64-bit mode Bits 4-1: Reserved Bit 5: LZCNT available Bits 7-6 Reserved Bit 8: PREFETCHW Bits 31-9: Reserved
	EDX	Bits 10-0: Reserved Bit 11: SYSCALL/SYSRET available (when in 64-bit mode) Bits 19-12: Reserved = 0 Bit 20: Execute Disable Bit available Bits 25-21: Reserved = 0 Bit 26: 1-GByte pages are available if 1 Bit 27: RDTSCP and IA32_TSC_AUX are available if 1 Bits 28: Reserved = 0 Bit 29: Intel® 64 Architecture available if 1 Bits 31-30: Reserved = 0

Table 2-8. Information Returned by CPUID Instruction(Continued)

Initial EAX Value	Information Provided about the Processor	
80000002H	EAX EBX ECX EDX	Processor Brand String Processor Brand String Continued Processor Brand String Continued Processor Brand String Continued
80000003H	EAX EBX ECX EDX	Processor Brand String Continued Processor Brand String Continued Processor Brand String Continued Processor Brand String Continued
80000004H	EAX EBX ECX EDX	Processor Brand String Continued Processor Brand String Continued Processor Brand String Continued Processor Brand String Continued
80000005H	EAX EBX ECX EDX	Reserved = 0 Reserved = 0 Reserved = 0 Reserved = 0
80000006H	EAX EBX ECX EDX	Reserved = 0 Reserved = 0 Bits 7-0: Cache Line size in bytes Bits 11-08: Reserved Bits 15-12: L2 Associativity field * Bits 31-16: Cache size in 1K units Reserved = 0 NOTES: * L2 associativity field encodings: 00H - Disabled 01H - Direct mapped 02H - 2-way 04H - 4-way 06H - 8-way 08H - 16-way 0FH - Fully associative
80000007H	EAX EBX ECX EDX	Reserved = 0 Reserved = 0 Reserved = 0 Bits 07-00: Reserved = 0 Bit 08: Invariant TSC available if 1 Bits 31-09: Reserved = 0
80000008H	EAX EBX ECX EDX	Virtual/Physical Address size Bits 7-0: #Physical Address Bits* Bits 15-8: #Virtual Address Bits Bits 31-16: Reserved = 0 Reserved = 0 Reserved = 0 Reserved = 0 NOTES: * If CPUID.80000008H:EAX[7:0] is supported, the maximum physical address number supported should come from this field.

INPUT EAX = 0H: Returns CPUID's Highest Value for Basic Processor Information and the Vendor Identification String

When CPUID executes with EAX set to 0H, the processor returns the highest value the CPUID recognizes for returning basic processor information. The value is returned in the EAX register (see Table 2-9) and is processor specific.

A vendor identification string is also returned in EBX, EDX, and ECX. For Intel processors, the string is "GenuineIntel" and is expressed:

EBX ← 756e6547h (* "Genu", with G in the low 4 bits of BL *)

EDX ← 49656e69h (* "inel", with i in the low 4 bits of DL *)

ECX ← 6c65746eh (* "ntel", with n in the low 4 bits of CL *)

INPUT EAX = 80000000H: Returns CPUID's Highest Value for Extended Processor Information

When CPUID executes with EAX set to 0H, the processor returns the highest value the processor recognizes for returning extended processor information. The value is returned in the EAX register (see Table 2-9) and is processor specific.

Table 2-9. Highest CPUID Source Operand for Intel 64 and IA-32 Processors

Intel 64 or IA-32 Processors	Highest Value in EAX	
	Basic Information	Extended Function Information
Earlier Intel486 Processors	CPUID Not Implemented	CPUID Not Implemented
Later Intel486 Processors and Pentium Processors	01H	Not Implemented
Pentium Pro and Pentium II Processors, Intel® Celeron® Processors	02H	Not Implemented
Pentium III Processors	03H	Not Implemented
Pentium 4 Processors	02H	80000004H
Intel Xeon Processors	02H	80000004H
Pentium M Processor	02H	80000004H
Pentium 4 Processor supporting Hyper-Threading Technology	05H	80000008H
Pentium D Processor (8xx)	05H	80000008H
Pentium D Processor (9xx)	06H	80000008H
Intel Core Duo Processor	0AH	80000008H
Intel Core 2 Duo Processor	0AH	80000008H
Intel Xeon Processor 3000, 5100, 5300 Series	0AH	80000008H
Intel Xeon Processor 3000, 5100, 5200, 5300, 5400 Series	0AH	80000008H
Intel Core 2 Duo Processor 8000 Series	0DH	80000008H
Intel Xeon Processor 5200, 5400 Series	0AH	80000008H

IA32_BIOS_SIGN_ID Returns Microcode Update Signature

For processors that support the microcode update facility, the IA32_BIOS_SIGN_ID MSR is loaded with the update signature whenever CPUID executes. The signature is returned in the upper DWORD. For details, see Chapter 10 in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*.

INPUT EAX = 01H: Returns Model, Family, Stepping Information

When CPUID executes with EAX set to 01H, version information is returned in EAX (see Figure 2-4). For example: model, family, and processor type for the Intel Xeon processor 5100 series is as follows:

- Model — 1111B
- Family — 0101B
- Processor Type — 00B

See Table 2-10 for available processor type values. Stepping IDs are provided as needed.

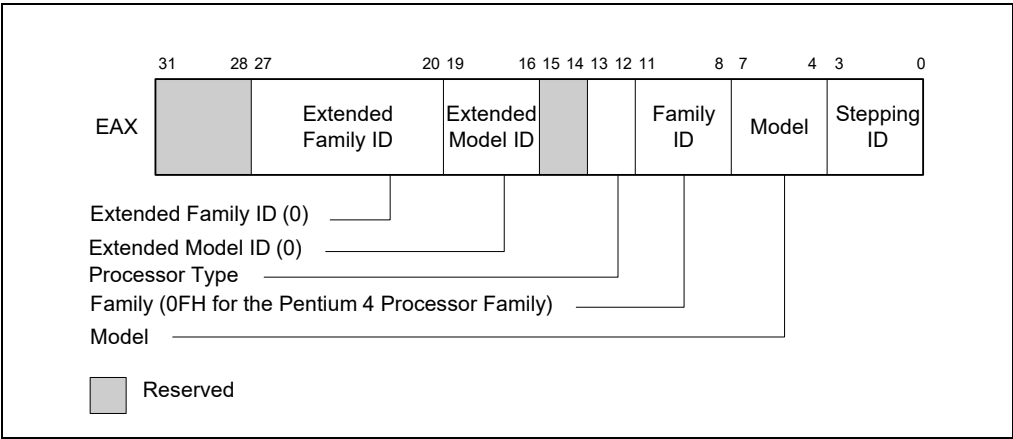


Figure 2-4. Version Information Returned by CPUID in EAX

Table 2-10. Processor Type Field

Type	Encoding
Original OEM Processor	00B
Intel OverDrive™ Processor	01B
Dual processor (not applicable to Intel486 processors)	10B
Intel reserved	11B

NOTE

See "Caching Translation Information" in Chapter 4, "Paging," in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A*, and Chapter 16 in the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1*, for information on identifying earlier IA-32 processors.

The Extended Family ID needs to be examined only when the Family ID is 0FH. Integrate the fields into a display using the following rule:

```

IF Family_ID ≠ 0FH
    THEN Displayed_Family = Family_ID;
    ELSE Displayed_Family = Extended_Family_ID + Family_ID;
    (* Right justify and zero-extend 4-bit field. *)
FI;
(* Show Display_Family as HEX field. *)
    
```

The Extended Model ID needs to be examined only when the Family ID is 06H or 0FH. Integrate the field into a display using the following rule:

```

IF (Family_ID = 06H or Family_ID = 0FH)
    THEN Displayed_Model = (Extended_Model_ID << 4) + Model_ID;
    
```

(* Right justify and zero-extend 4-bit field; display Model_ID as HEX field.*)

ELSE Displayed_Model = Model_ID;

FI;

(* Show Display_Model as HEX field. *)

INPUT EAX = 01H: Returns Additional Information in EBX

When CPUID executes with EAX set to 01H, additional information is returned to the EBX register:

- Brand index (low byte of EBX) — this number provides an entry into a brand string table that contains brand strings for IA-32 processors. More information about this field is provided later in this section.
- CLFLUSH instruction cache line size (second byte of EBX) — this number indicates the size of the cache line flushed with CLFLUSH instruction in 8-byte increments. This field was introduced in the Pentium 4 processor.
- Local APIC ID (high byte of EBX) — this number is the 8-bit ID that is assigned to the local APIC on the processor during power up. This field was introduced in the Pentium 4 processor.

INPUT EAX = 01H: Returns Feature Information in ECX and EDX

When CPUID executes with EAX set to 01H, feature information is returned in ECX and EDX.

- Figure 2-5 and Table 2-11 show encodings for ECX.
- Figure 2-6 and Table 2-12 show encodings for EDX.

For all feature flags, a 1 indicates that the feature is supported. Use Intel to properly interpret feature flags.

NOTE

Software must confirm that a processor feature is present using feature flags returned by CPUID prior to using the feature. Software should not depend on future offerings retaining all features.

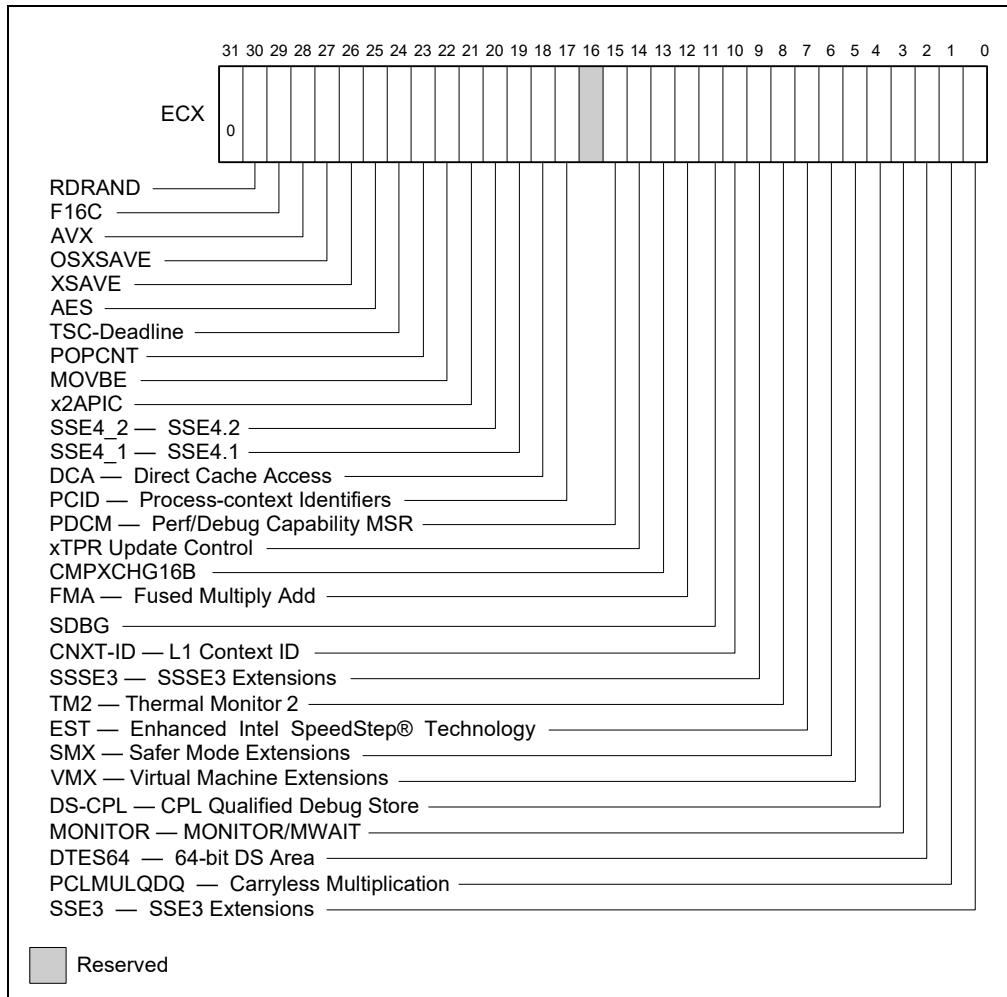


Figure 2-5. Feature Information Returned in the ECX Register

Table 2-11. Feature Information Returned in the ECX Register

Bit #	Mnemonic	Description
0	SSE3	Streaming SIMD Extensions 3 (SSE3). A value of 1 indicates the processor supports this technology.
1	PCLMULQDQ	A value of 1 indicates the processor supports PCLMULQDQ instruction.
2	DTES64	64-bit DS Area. A value of 1 indicates the processor supports DS area using 64-bit layout.
3	MONITOR	MONITOR/MWAIT. A value of 1 indicates the processor supports this feature.
4	DS-CPL	CPL Qualified Debug Store. A value of 1 indicates the processor supports the extensions to the Debug Store feature to allow for branch message storage qualified by CPL.
5	VMX	Virtual Machine Extensions. A value of 1 indicates that the processor supports this technology.
6	SMX	Safer Mode Extensions. A value of 1 indicates that the processor supports this technology. See Chapter 5, “Safer Mode Extensions Reference”.
7	EST	Enhanced Intel SpeedStep® Technology. A value of 1 indicates that the processor supports this technology.
8	TM2	Thermal Monitor 2. A value of 1 indicates whether the processor supports this technology.
9	SSSE3	A value of 1 indicates the presence of the Supplemental Streaming SIMD Extensions 3 (SSSE3). A value of 0 indicates the instruction extensions are not present in the processor.

Table 2-11. Feature Information Returned in the ECX Register (Continued)

Bit #	Mnemonic	Description
10	CNXT-ID	L1 Context ID. A value of 1 indicates the L1 data cache mode can be set to either adaptive mode or shared mode. A value of 0 indicates this feature is not supported. See definition of the IA32_MISC_ENABLE MSR Bit 24 (L1 Data Cache Context Mode) for details.
11	SDBG	A value of 1 indicates the processor supports IA32_DEBUG_INTERFACE MSR for silicon debug.
12	FMA	A value of 1 indicates the processor supports FMA extensions using YMM state.
13	CMPXCHG16B	CMPXCHG16B Available. A value of 1 indicates that the feature is available.
14	xTPR Update Control	xTPR Update Control. A value of 1 indicates that the processor supports changing IA32_MISC_ENABLE[bit 23].
15	PDCM	Perfmon and Debug Capability. A value of 1 indicates the processor supports the performance and debug feature indication MSR IA32_PERF_CAPABILITIES.
16	Reserved	Reserved
17	PCID	Process-context identifiers. A value of 1 indicates that the processor supports PCIDs and that software may set CR4.PCIDE to 1.
18	DCA	A value of 1 indicates the processor supports the ability to prefetch data from a memory mapped device.
19	SSE4.1	A value of 1 indicates that the processor supports SSE4.1.
20	SSE4.2	A value of 1 indicates that the processor supports SSE4.2.
21	x2APIC	A value of 1 indicates that the processor supports x2APIC feature.
22	MOVBE	A value of 1 indicates that the processor supports MOVBE instruction.
23	POPCNT	A value of 1 indicates that the processor supports the POPCNT instruction.
24	TSC-Deadline	A value of 1 indicates that the processor's local APIC timer supports one-shot operation using a TSC deadline value.
25	AES	A value of 1 indicates that the processor supports the AESNI instruction extensions.
26	XSAVE	A value of 1 indicates that the processor supports the XSAVE/XRSTOR processor extended states feature, the XSETBV/XGETBV instructions, and XCRO.
27	OSXSAVE	A value of 1 indicates that the OS has set CR4.OSXSAVE[bit 18] to enable XSETBV/XGETBV instructions to access XCRO and to support processor extended state management using XSAVE/XRSTOR.
28	AVX	A value of 1 indicates that processor supports AVX instructions operating on 256-bit YMM state, and three-operand encoding of 256-bit and 128-bit SIMD instructions.
29	F16C	A value of 1 indicates that processor supports 16-bit floating-point conversion instructions.
30	RDRAND	A value of 1 indicates that processor supports RDRAND instruction.
31	Not Used	Always return 0.

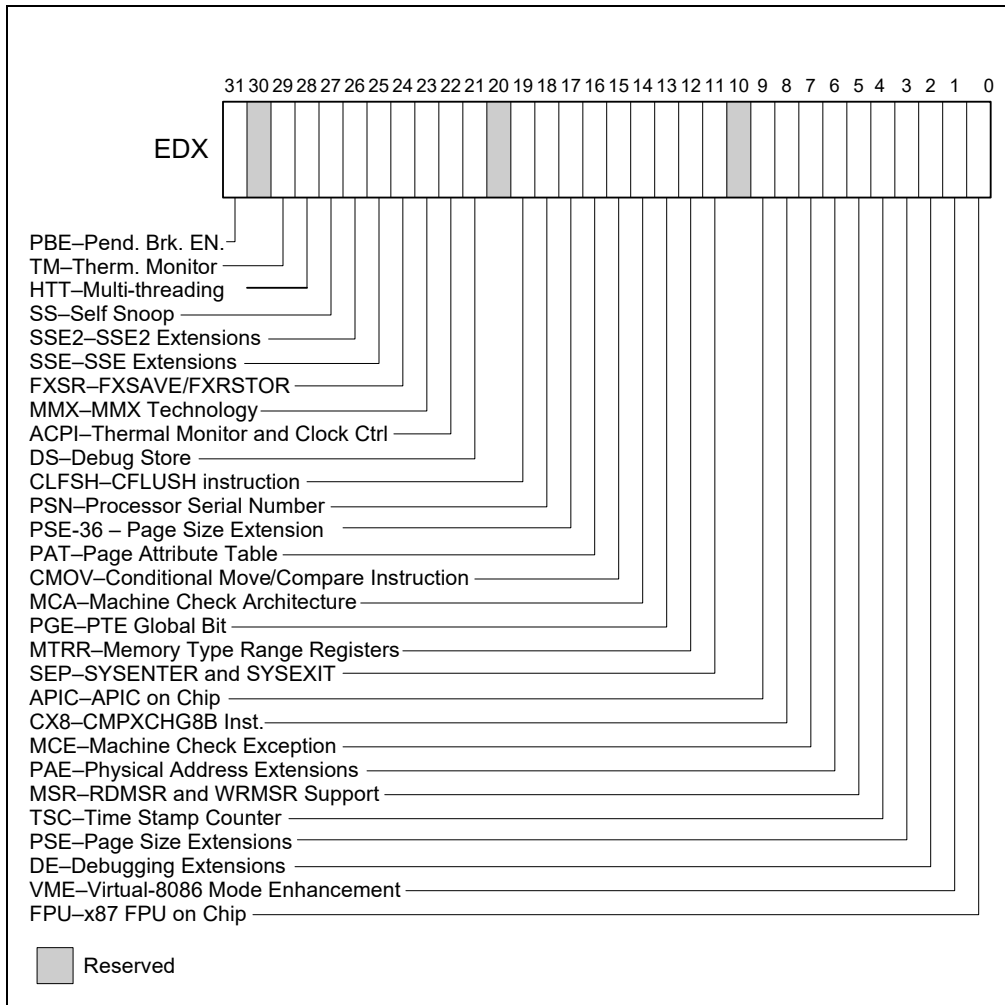


Figure 2-6. Feature Information Returned in the EDX Register

Table 2-12. More on Feature Information Returned in the EDX Register

Bit #	Mnemonic	Description
0	FPU	Floating-point Unit On-Chip. The processor contains an x87 FPU.
1	VME	Virtual 8086 Mode Enhancements. Virtual 8086 mode enhancements, including CR4.VME for controlling the feature, CR4.PVI for protected mode virtual interrupts, software interrupt indirection, expansion of the TSS with the software indirection bitmap, and EFLAGS.VIF and EFLAGS.VIP flags.
2	DE	Debugging Extensions. Support for I/O breakpoints, including CR4.DE for controlling the feature, and optional trapping of accesses to DR4 and DR5.
3	PSE	Page Size Extension. Large pages of size 4 MByte are supported, including CR4.PSE for controlling the feature, the defined dirty bit in PDE (Page Directory Entries), optional reserved bit trapping in CR3, PDEs, and PTEs.
4	TSC	Time Stamp Counter. The RDTSC instruction is supported, including CR4.TSD for controlling privilege.
5	MSR	Model Specific Registers RDMSR and WRMSR Instructions. The RDMSR and WRMSR instructions are supported. Some of the MSRs are implementation dependent.

Table 2-12. More on Feature Information Returned in the EDX Register(Continued)

Bit #	Mnemonic	Description
6	PAE	Physical Address Extension. Physical addresses greater than 32 bits are supported: extended page table entry formats, an extra level in the page translation tables is defined, 2-MByte pages are supported instead of 4 Mbyte pages if PAE bit is 1. The actual number of address bits beyond 32 is not defined, and is implementation specific.
7	MCE	Machine Check Exception. Exception 18 is defined for Machine Checks, including CR4.MCE for controlling the feature. This feature does not define the model-specific implementations of machine-check error logging, reporting, and processor shutdowns. Machine Check exception handlers may have to depend on processor version to do model specific processing of the exception, or test for the presence of the Machine Check feature.
8	CX8	CMPXCHG8B Instruction. The compare-and-exchange 8 bytes (64 bits) instruction is supported (implicitly locked and atomic).
9	APIC	APIC On-Chip. The processor contains an Advanced Programmable Interrupt Controller (APIC), responding to memory mapped commands in the physical address range FFFE0000H to FFFE0FFFH (by default - some processors permit the APIC to be relocated).
10	Reserved	Reserved
11	SEP	SYSENTER and SYSEXIT Instructions. The SYSENTER and SYSEXIT and associated MSRs are supported.
12	MTRR	Memory Type Range Registers. MTRRs are supported. The MTRRcap MSR contains feature bits that describe what memory types are supported, how many variable MTRRs are supported, and whether fixed MTRRs are supported.
13	PGE	Page Global Bit. The global bit is supported in paging-structure entries that map a page, indicating TLB entries that are common to different processes and need not be flushed. The CR4.PGE bit controls this feature.
14	MCA	Machine Check Architecture. The Machine Check Architecture, which provides a compatible mechanism for error reporting in P6 family, Pentium 4, Intel Xeon processors, and future processors, is supported. The MCG_CAP MSR contains feature bits describing how many banks of error reporting MSRs are supported.
15	CMOV	Conditional Move Instructions. The conditional move instruction CMOV is supported. In addition, if x87 FPU is present as indicated by the CPUID.FPU feature bit, then the FCOMI and FCMOV instructions are supported
16	PAT	Page Attribute Table. Page Attribute Table is supported. This feature augments the Memory Type Range Registers (MTRRs), allowing an operating system to specify attributes of memory accessed through a linear address on a 4KB granularity.
17	PSE-36	36-Bit Page Size Extension. 4-MByte pages addressing physical memory beyond 4 GBytes are supported with 32-bit paging. This feature indicates that upper bits of the physical address of a 4-MByte page are encoded in bits 20:13 of the page-directory entry. Such physical addresses are limited by MAXPHYADDR and may be up to 40 bits in size.
18	PSN	Processor Serial Number. The processor supports the 96-bit processor identification number feature and the feature is enabled.
19	CLFSH	CLFLUSH Instruction. CLFLUSH Instruction is supported.
20	Reserved	Reserved
21	DS	Debug Store. The processor supports the ability to write debug information into a memory resident buffer. This feature is used by the branch trace store (BTS) and precise event-based sampling (PEBS) facilities (see Chapter 23, "Introduction to Virtual-Machine Extensions," in the <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3C</i>).
22	ACPI	Thermal Monitor and Software Controlled Clock Facilities. The processor implements internal MSRs that allow processor temperature to be monitored and processor performance to be modulated in predefined duty cycles under software control.
23	MMX	Intel MMX Technology. The processor supports the Intel MMX technology.
24	FXSR	FXSAVE and FXRSTOR Instructions. The FXSAVE and FXRSTOR instructions are supported for fast save and restore of the floating-point context. Presence of this bit also indicates that CR4.OSFXSR is available for an operating system to indicate that it supports the FXSAVE and FXRSTOR instructions.

Table 2-12. More on Feature Information Returned in the EDX Register(Continued)

Bit #	Mnemonic	Description
25	SSE	SSE. The processor supports the SSE extensions.
26	SSE2	SSE2. The processor supports the SSE2 extensions.
27	SS	Self Snoop. The processor supports the management of conflicting memory types by performing a snoop of its own cache structure for transactions issued to the bus.
28	HTT	Max APIC IDs reserved field is Valid. A value of 0 for HTT indicates there is only a single logical processor in the package and software should assume only a single APIC ID is reserved. A value of 1 for HTT indicates the value in CPUID.1.EBX[23:16] (the Maximum number of addressable IDs for logical processors in this package) is valid for the package.
29	TM	Thermal Monitor. The processor implements the thermal monitor automatic thermal control circuitry (TCC).
30	Reserved	Reserved
31	PBE	Pending Break Enable. The processor supports the use of the FERR#/PBE# pin when the processor is in the stop-clock state (STPCLK# is asserted) to signal the processor that an interrupt is pending and that the processor should return to normal operation to handle the interrupt. Bit 10 (PBE enable) in the IA32_MISC_ENABLE MSR enables this capability.

INPUT EAX = 02H: Cache and TLB Information Returned in EAX, EBX, ECX, EDX

When CPUID executes with EAX set to 02H, the processor returns information about the processor's internal caches and TLBs in the EAX, EBX, ECX, and EDX registers.

The encoding is as follows:

- The least-significant byte in register EAX (register AL) indicates the number of times the CPUID instruction must be executed with an input value of 02H to get a complete description of the processor's caches and TLBs. The first member of the family of Pentium 4 processors will return a 01H.
- The most significant bit (bit 31) of each register indicates whether the register contains valid information (set to 0) or is reserved (set to 1).
- If a register contains valid information, the information is contained in 1 byte descriptors. Table 2-13 shows the encoding of these descriptors. Note that the order of descriptors in the EAX, EBX, ECX, and EDX registers is not defined; that is, specific bytes are not designated to contain descriptors for specific cache or TLB types. The descriptors may appear in any order.

Table 2-13. Encoding of Cache and TLB Descriptors

Descriptor Value	Cache or TLB Description
00H	Null descriptor
01H	Instruction TLB: 4 KByte pages, 4-way set associative, 32 entries
02H	Instruction TLB: 4 MByte pages, 4-way set associative, 2 entries
03H	Data TLB: 4 KByte pages, 4-way set associative, 64 entries
04H	Data TLB: 4 MByte pages, 4-way set associative, 8 entries
05H	Data TLB1: 4 MByte pages, 4-way set associative, 32 entries
06H	1st-level instruction cache: 8 KBytes, 4-way set associative, 32 byte line size
08H	1st-level instruction cache: 16 KBytes, 4-way set associative, 32 byte line size
0AH	1st-level data cache: 8 KBytes, 2-way set associative, 32 byte line size
0BH	Instruction TLB: 4 MByte pages, 4-way set associative, 4 entries
0CH	1st-level data cache: 16 KBytes, 4-way set associative, 32 byte line size
22H	3rd-level cache: 512 KBytes, 4-way set associative, 64 byte line size, 2 lines per sector
23H	3rd-level cache: 1 MBytes, 8-way set associative, 64 byte line size, 2 lines per sector
25H	3rd-level cache: 2 MBytes, 8-way set associative, 64 byte line size, 2 lines per sector

Table 2-13. Encoding of Cache and TLB Descriptors (Continued)

Descriptor Value	Cache or TLB Description
29H	3rd-level cache: 4 MBytes, 8-way set associative, 64 byte line size, 2 lines per sector
2CH	1st-level data cache: 32 KBytes, 8-way set associative, 64 byte line size
30H	1st-level instruction cache: 32 KBytes, 8-way set associative, 64 byte line size
40H	No 2nd-level cache or, if processor contains a valid 2nd-level cache, no 3rd-level cache
41H	2nd-level cache: 128 KBytes, 4-way set associative, 32 byte line size
42H	2nd-level cache: 256 KBytes, 4-way set associative, 32 byte line size
43H	2nd-level cache: 512 KBytes, 4-way set associative, 32 byte line size
44H	2nd-level cache: 1 MByte, 4-way set associative, 32 byte line size
45H	2nd-level cache: 2 MByte, 4-way set associative, 32 byte line size
46H	3rd-level cache: 4 MByte, 4-way set associative, 64 byte line size
47H	3rd-level cache: 8 MByte, 8-way set associative, 64 byte line size
49H	3rd-level cache: 4MB, 16-way set associative, 64-byte line size (Intel Xeon processor MP, Family 0FH, Model 06H); 2nd-level cache: 4 MByte, 16-way set associative, 64 byte line size
4AH	3rd-level cache: 6MByte, 12-way set associative, 64 byte line size
4BH	3rd-level cache: 8MByte, 16-way set associative, 64 byte line size
4CH	3rd-level cache: 12MByte, 12-way set associative, 64 byte line size
4DH	3rd-level cache: 16MByte, 16-way set associative, 64 byte line size
4EH	2nd-level cache: 6MByte, 24-way set associative, 64 byte line size
50H	Instruction TLB: 4 KByte and 2-MByte or 4-MByte pages, 64 entries
51H	Instruction TLB: 4 KByte and 2-MByte or 4-MByte pages, 128 entries
52H	Instruction TLB: 4 KByte and 2-MByte or 4-MByte pages, 256 entries
56H	Data TLB0: 4 MByte pages, 4-way set associative, 16 entries
57H	Data TLB0: 4 KByte pages, 4-way associative, 16 entries
5BH	Data TLB: 4 KByte and 4 MByte pages, 64 entries
5CH	Data TLB: 4 KByte and 4 MByte pages, 128 entries
5DH	Data TLB: 4 KByte and 4 MByte pages, 256 entries
60H	1st-level data cache: 16 KByte, 8-way set associative, 64 byte line size
66H	1st-level data cache: 8 KByte, 4-way set associative, 64 byte line size
67H	1st-level data cache: 16 KByte, 4-way set associative, 64 byte line size
68H	1st-level data cache: 32 KByte, 4-way set associative, 64 byte line size
70H	Trace cache: 12 K-μop, 8-way set associative
71H	Trace cache: 16 K-μop, 8-way set associative
72H	Trace cache: 32 K-μop, 8-way set associative
78H	2nd-level cache: 1 MByte, 4-way set associative, 64byte line size
79H	2nd-level cache: 128 KByte, 8-way set associative, 64 byte line size, 2 lines per sector
7AH	2nd-level cache: 256 KByte, 8-way set associative, 64 byte line size, 2 lines per sector
7BH	2nd-level cache: 512 KByte, 8-way set associative, 64 byte line size, 2 lines per sector
7CH	2nd-level cache: 1 MByte, 8-way set associative, 64 byte line size, 2 lines per sector
7DH	2nd-level cache: 2 MByte, 8-way set associative, 64byte line size
7FH	2nd-level cache: 512 KByte, 2-way set associative, 64-byte line size
82H	2nd-level cache: 256 KByte, 8-way set associative, 32 byte line size

Table 2-13. Encoding of Cache and TLB Descriptors (Continued)

Descriptor Value	Cache or TLB Description
83H	2nd-level cache: 512 KByte, 8-way set associative, 32 byte line size
84H	2nd-level cache: 1 MByte, 8-way set associative, 32 byte line size
85H	2nd-level cache: 2 MByte, 8-way set associative, 32 byte line size
86H	2nd-level cache: 512 KByte, 4-way set associative, 64 byte line size
87H	2nd-level cache: 1 MByte, 8-way set associative, 64 byte line size
B0H	Instruction TLB: 4 KByte pages, 4-way set associative, 128 entries
B1H	Instruction TLB: 2M pages, 4-way, 8 entries or 4M pages, 4-way, 4 entries
B3H	Data TLB: 4 KByte pages, 4-way set associative, 128 entries
B4H	Data TLB1: 4 KByte pages, 4-way associative, 256 entries
F0H	64-Byte prefetching
F1H	128-Byte prefetching

Example 2-1. Example of Cache and TLB Interpretation

The first member of the family of Pentium 4 processors returns the following information about caches and TLBs when the CPUID executes with an input value of 2:

```

EAX    66 5B 50 01H
EBX     0H
ECX     0H
EDX    00 7A 70 00H

```

Which means:

- The least-significant byte (byte 0) of register EAX is set to 01H. This indicates that CPUID needs to be executed once with an input value of 2 to retrieve complete information about caches and TLBs.
- The most-significant bit of all four registers (EAX, EBX, ECX, and EDX) is set to 0, indicating that each register contains valid 1-byte descriptors.
- Bytes 1, 2, and 3 of register EAX indicate that the processor has:
 - 50H - a 64-entry instruction TLB, for mapping 4-KByte and 2-MByte or 4-MByte pages.
 - 5BH - a 64-entry data TLB, for mapping 4-KByte and 4-MByte pages.
 - 66H - an 8-KByte 1st level data cache, 4-way set associative, with a 64-Byte cache line size.
- The descriptors in registers EBX and ECX are valid, but contain NULL descriptors.
- Bytes 0, 1, 2, and 3 of register EDX indicate that the processor has:
 - 00H - NULL descriptor.
 - 70H - Trace cache: 12 K- μ op, 8-way set associative.
 - 7AH - a 256-KByte 2nd level cache, 8-way set associative, with a sectored, 64-byte cache line size.
 - 00H - NULL descriptor.

INPUT EAX = 04H: Returns Deterministic Cache Parameters for Each Level

When CPUID executes with EAX set to 04H and ECX contains an index value, the processor returns encoded data that describe a set of deterministic cache parameters (for the cache level associated with the input in ECX). Valid index values start from 0.

Software can enumerate the deterministic cache parameters for each level of the cache hierarchy starting with an index value of 0, until the parameters report the value associated with the cache type field is 0. The architecturally defined fields reported by deterministic cache parameters are documented in Table 2-8.

The CPUID leaf 4 also reports data that can be used to derive the topology of processor cores in a physical package. This information is constant for all valid index values. Software can query the raw data reported by executing CPUID with EAX=04H and ECX=0H and use it as part of the topology enumeration algorithm described in Chapter 8, “Multiple-Processor Management,” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*.

INPUT EAX = 05H: Returns MONITOR and MWAIT Features

When CPUID executes with EAX set to 05H, the processor returns information about features available to MONITOR/MWAIT instructions. The MONITOR instruction is used for address-range monitoring in conjunction with MWAIT instruction. The MWAIT instruction optionally provides additional extensions for advanced power management. See Table 2-8.

INPUT EAX = 06H: Returns Thermal and Power Management Features

When CPUID executes with EAX set to 06H, the processor returns information about thermal and power management features. See Table 2-8.

INPUT EAX = 07H: Returns Structured Extended Feature Enumeration Information

When CPUID executes with EAX set to 07H and ECX = 0H, the processor returns information about the maximum number of sub-leaves that contain extended feature flags. See Table 2-8.

When CPUID executes with EAX set to 07H and ECX = n (n > 1 and less than the number of non-zero bits in CPUID.(EAX=07H, ECX= 0H).EAX, the processor returns information about extended feature flags. See Table 2-8. In sub-leaf 0, only EAX has the number of sub-leaves. In sub-leaf 0, EBX, ECX & EDX all contain extended feature flags.

Table 2-14. Structured Extended Feature Leaf, Function 0, EBX Register

Bit #	Mnemonic	Description
0	RWFSGBASE	A value of 1 indicates the processor supports RD/WR FSGSBASE instructions
1-31	Reserved	Reserved

INPUT EAX = 09H: Returns Direct Cache Access Information

When CPUID executes with EAX set to 09H, the processor returns information about Direct Cache Access capabilities. See Table 2-8.

INPUT EAX = 0AH: Returns Architectural Performance Monitoring Features

When CPUID executes with EAX set to 0AH, the processor returns information about support for architectural performance monitoring capabilities. Architectural performance monitoring is supported if the version ID (see Table 2-8) is greater than Pn 0. See Table 2-8.

For each version of architectural performance monitoring capability, software must enumerate this leaf to discover the programming facilities and the architectural performance events available in the processor. The details are described in Chapter 17, “Debug, Branch Profile, TSC, and Quality of Service,” in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A*.

INPUT EAX = 0BH: Returns Extended Topology Information

When CPUID executes with EAX set to 0BH, the processor returns information about extended topology enumeration data. Software must detect the presence of CPUID leaf 0BH by verifying (a) the highest leaf index supported by CPUID is >= 0BH, and (b) CPUID.0BH:EBX[15:0] reports a non-zero value. See Table 2-8.

INPUT EAX = 0DH: Returns Processor Extended States Enumeration Information

When CPUID executes with EAX set to 0DH and ECX = 0H, the processor returns information about the bit-vector representation of all processor state extensions that are supported in the processor and storage size requirements of the XSAVE/XRSTOR area. See Table 2-8.

When CPUID executes with EAX set to 0DH and ECX = n (n > 1, and is a valid sub-leaf index), the processor returns information about the size and offset of each processor extended state save area within the XSAVE/XRSTOR area. See Table 2-8. Software can use the forward-extendable technique depicted below to query the valid sub-leaves and obtain size and offset information for each processor extended state save area:

```
For i = 2 to 62 // sub-leaf 1 is reserved
  IF (CPUID.(EAX=0DH, ECX=0):VECTOR[i] = 1) // VECTOR is the 64-bit value of EDX:EAX
    Execute CPUID.(EAX=0DH, ECX = i) to examine size and offset for sub-leaf i;
  FI;
```

INPUT EAX = 0FH: Returns Platform Quality of Service (PQoS) Monitoring Enumeration Information

When CPUID executes with EAX set to 0FH and ECX = 0, the processor returns information about the bit-vector representation of QoS monitoring resource types that are supported in the processor and maximum range of RMID values the processor can use to monitor of any supported resource types. Each bit, starting from bit 1, corresponds to a specific resource type if the bit is set. The bit position corresponds to the sub-leaf index (or ResID) that software must use to query QoS monitoring capability available for that type. See Table 2-8.

When CPUID executes with EAX set to 0FH and ECX = n (n >= 1, and is a valid ResID), the processor returns information software can use to program IA32_PQR_ASSOC, IA32_QM_EVTSEL MSRs before reading QoS data from the IA32_QM_CTR MSR.

INPUT EAX = 10H: Returns Platform Quality of Service (PQoS) Enforcement Enumeration Information

When CPUID executes with EAX set to 10H and ECX = 0, the processor returns information about the bit-vector representation of QoS Enforcement resource types that are supported in the processor. Each bit, starting from bit 1, corresponds to a specific resource type if the bit is set. The bit position corresponds to the sub-leaf index (or ResID) that software must use to query QoS enforcement capability available for that type. See Table 2-8.

When CPUID executes with EAX set to 10H and ECX = n (n >= 1, and is a valid ResID), the processor returns information about available classes of service and range of QoS mask MSRs that software can use to configure each class of services using capability bit masks in the QoS Mask registers, IA32_resourceType_Mask_n.

INPUT EAX = 14H: Returns Intel Processor Trace Enumeration Information

When CPUID executes with EAX set to 14H and ECX = 0H, the processor returns information about Intel Processor Trace extensions. See Table 2-8.

When CPUID executes with EAX set to 14H and ECX = n (n > 0 and less than the number of non-zero bits in CPUID.(EAX=14H, ECX= 0H).EAX), the processor returns information about packet generation in Intel Processor Trace. See Table 2-8.

INPUT EAX = 15H: Returns Time Stamp Counter and Core Crystal Clock Information

When CPUID executes with EAX set to 15H and ECX = 0H, the processor returns information about Time Stamp Counter and Core Crystal Clock. See Table 2-8.

INPUT EAX = 16H: Returns Processor Frequency Information

When CPUID executes with EAX set to 16H, the processor returns information about Processor Frequency Information. See Table 2-8.

INPUT EAX = 17H: Returns System-On-Chip Information

When CPUID executes with EAX set to 17H, the processor returns information about the System-On-Chip Vendor Attribute Enumeration. See Table 2-8.

INPUT EAX = 18H: Returns Deterministic Address Translation Parameters Information

When CUID executes with EAX set to 18H, the processor returns information about the Deterministic Address Translation Parameters. See Table 2-8.

METHODS FOR RETURNING BRANDING INFORMATION

Use the following techniques to access branding information:

- 1. Processor brand string method; this method also returns the processor’s maximum operating frequency
- 2. Processor brand index; this method uses a software supplied brand string table.

These two methods are discussed in the following sections. For methods that are available in early processors, see Section: “Identification of Earlier IA-32 Processors” in Chapter 16 of the *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 1*.

The Processor Brand String Method

Figure 2-7 describes the algorithm used for detection of the brand string. Processor brand identification software should execute this algorithm on all Intel 64 and IA-32 processors.

This method (introduced with Pentium 4 processors) returns an ASCII brand identification string and the maximum operating frequency of the processor to the EAX, EBX, ECX, and EDX registers.

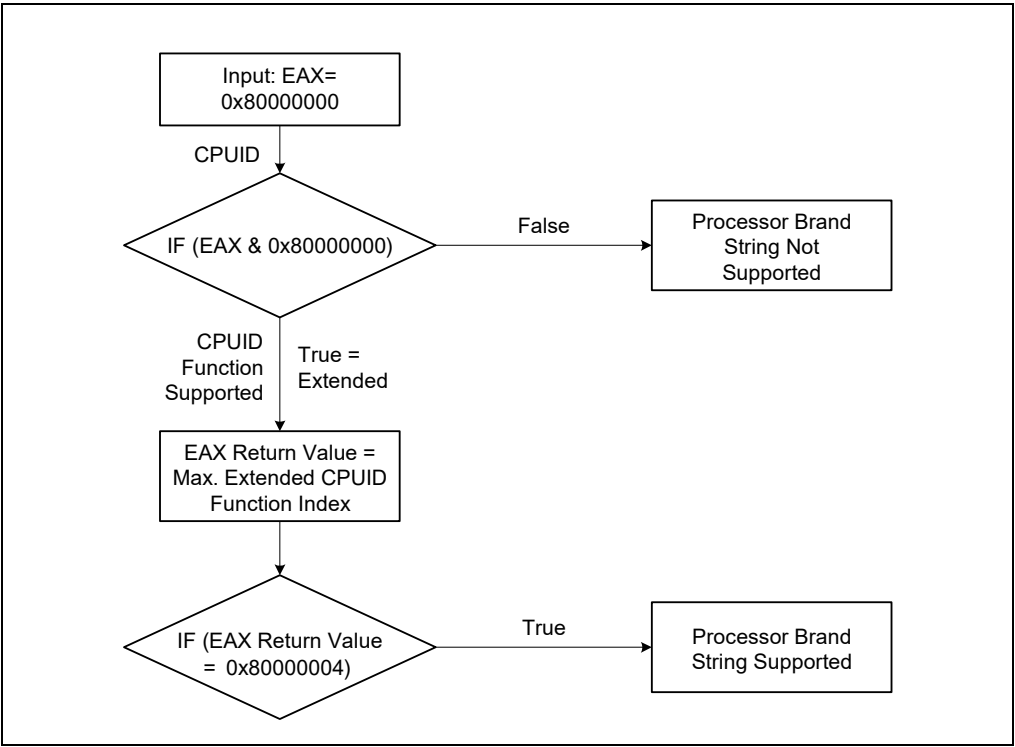


Figure 2-7. Determination of Support for the Processor Brand String

How Brand Strings Work

To use the brand string method, execute CUID with EAX input of 8000002H through 80000004H. For each input value, CUID returns 16 ASCII characters using EAX, EBX, ECX, and EDX. The returned string will be NULL-terminated.

Table 2-15 shows the brand string that is returned by the first processor in the Pentium 4 processor family.

Table 2-15. Processor Brand String Returned with Pentium 4 Processor

EAX Input Value	Return Values	ASCII Equivalent
80000002H	EAX = 20202020H EBX = 20202020H ECX = 20202020H EDX = 6E492020H	" " " " " " "nI "
80000003H	EAX = 286C6574H EBX = 50202952H ECX = 69746E65H EDX = 52286D75H	"(let" "P)R" "itne" "R(mu"
80000004H	EAX = 20342029H EBX = 20555043H ECX = 30303531H EDX = 007A484DH	" 4)" " UPC" "0051" "\0zHM"

Extracting the Maximum Processor Frequency from Brand Strings

Figure 2-8 provides an algorithm which software can use to extract the maximum processor operating frequency from the processor brand string.

NOTE

When a frequency is given in a brand string, it is the maximum qualified frequency of the processor, not the frequency at which the processor is currently running.

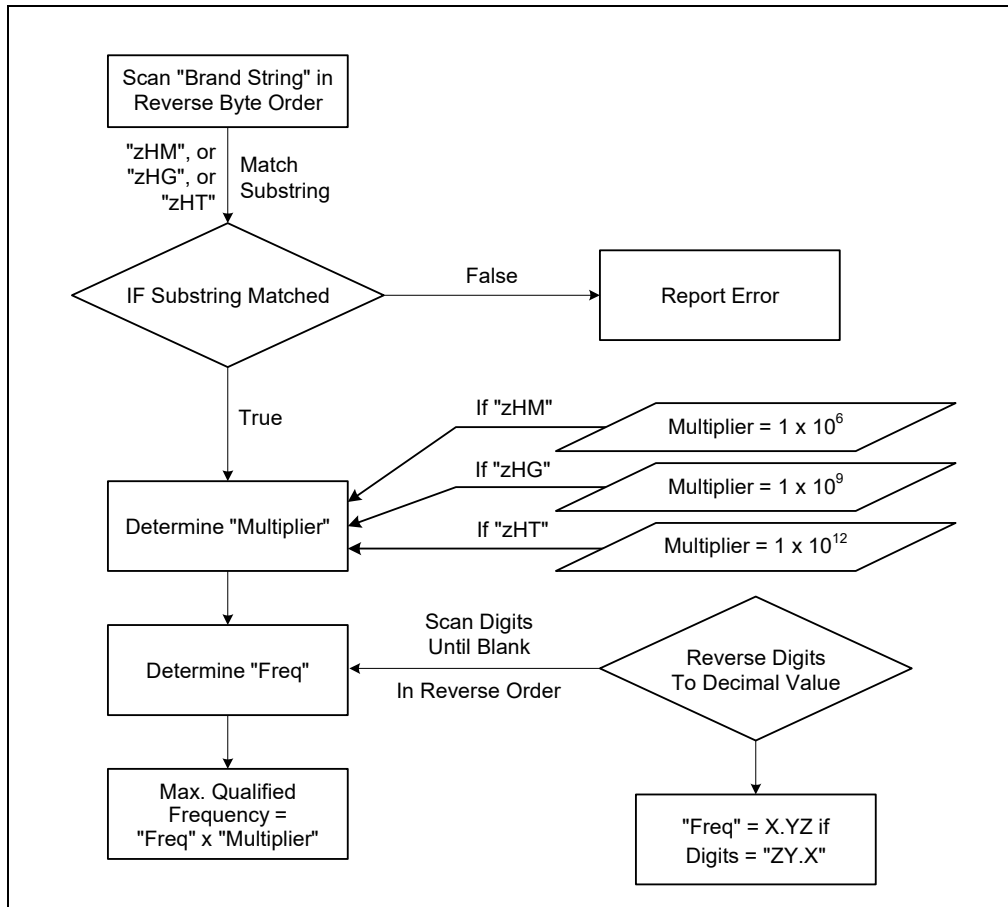


Figure 2-8. Algorithm for Extracting Maximum Processor Frequency

The Processor Brand Index Method

The brand index method (introduced with Pentium® III Xeon® processors) provides an entry point into a brand identification table that is maintained in memory by system software and is accessible from system- and user-level code. In this table, each brand index is associated with an ASCII brand identification string that identifies the official Intel family and model number of a processor.

When CPUID executes with EAX set to 01H, the processor returns a brand index to the low byte in EBX. Software can then use this index to locate the brand identification string for the processor in the brand identification table. The first entry (brand index 0) in this table is reserved, allowing for backward compatibility with processors that do not support the brand identification feature. Starting with processor signature family ID = 0FH, model = 03H, brand index method is no longer supported. Use brand string method instead.

Table 2-16 shows brand indices that have identification strings associated with them.

Table 2-16. Mapping of Brand Indices; and Intel 64 and IA-32 Processor Brand Strings

Brand Index	Brand String
00H	This processor does not support the brand identification feature
01H	Intel(R) Celeron(R) processor ¹
02H	Intel(R) Pentium(R) III processor ¹
03H	Intel(R) Pentium(R) III Xeon(R) processor; If processor signature = 000006B1h, then Intel(R) Celeron(R) processor
04H	Intel(R) Pentium(R) III processor
06H	Mobile Intel(R) Pentium(R) III processor-M
07H	Mobile Intel(R) Celeron(R) processor ¹
08H	Intel(R) Pentium(R) 4 processor
09H	Intel(R) Pentium(R) 4 processor
0AH	Intel(R) Celeron(R) processor ¹
0BH	Intel(R) Xeon(R) processor; If processor signature = 00000F13h, then Intel(R) Xeon(R) processor MP
0CH	Intel(R) Xeon(R) processor MP
0EH	Mobile Intel(R) Pentium(R) 4 processor-M; If processor signature = 00000F13h, then Intel(R) Xeon(R) processor
0FH	Mobile Intel(R) Celeron(R) processor ¹
11H	Mobile Genuine Intel(R) processor
12H	Intel(R) Celeron(R) M processor
13H	Mobile Intel(R) Celeron(R) processor ¹
14H	Intel(R) Celeron(R) processor
15H	Mobile Genuine Intel(R) processor
16H	Intel(R) Pentium(R) M processor
17H	Mobile Intel(R) Celeron(R) processor ¹
18H - 0FFH	RESERVED

NOTES:

1. Indicates versions of these processors that were introduced after the Pentium III

IA-32 Architecture Compatibility

CPUID is not supported in early models of the Intel486 processor or in any IA-32 processor earlier than the Intel486 processor.

Operation

IA32_BIOS_SIGN_ID MSR ← Update with installed microcode revision number;

CASE (EAX) OF

EAX = 0:

EAX ← Highest basic function input value understood by CPUID;

EBX ← Vendor identification string;

EDX ← Vendor identification string;

ECX ← Vendor identification string;

BREAK;

EAX = 1H:

EAX[3:0] ← Stepping ID;

```

EAX[7:4] ← Model;
EAX[11:8] ← Family;
EAX[13:12] ← Processor type;
EAX[15:14] ← Reserved;
EAX[19:16] ← Extended Model;
EAX[27:20] ← Extended Family;
EAX[31:28] ← Reserved;
EBX[7:0] ← Brand Index; (* Reserved if the value is zero. *)
EBX[15:8] ← CLFLUSH Line Size;
EBX[16:23] ← Reserved; (* Number of threads enabled = 2 if MT enable fuse set. *)
EBX[24:31] ← Initial APIC ID;
ECX ← Feature flags; (* See Figure 2-5. *)
EDX ← Feature flags; (* See Figure 2-6. *)
BREAK;
EAX = 2H:
    EAX ← Cache and TLB information;
    EBX ← Cache and TLB information;
    ECX ← Cache and TLB information;
    EDX ← Cache and TLB information;
BREAK;
EAX = 3H:
    EAX ← Reserved;
    EBX ← Reserved;
    ECX ← ProcessorSerialNumber[31:0];
    (* Pentium III processors only, otherwise reserved. *)
    EDX ← ProcessorSerialNumber[63:32];
    (* Pentium III processors only, otherwise reserved. *)
BREAK;
EAX = 4H:
    EAX ← Deterministic Cache Parameters Leaf; (* See Table 2-8. *)
    EBX ← Deterministic Cache Parameters Leaf;
    ECX ← Deterministic Cache Parameters Leaf;
    EDX ← Deterministic Cache Parameters Leaf;
BREAK;
EAX = 5H:
    EAX ← MONITOR/MWAIT Leaf; (* See Table 2-8. *)
    EBX ← MONITOR/MWAIT Leaf;
    ECX ← MONITOR/MWAIT Leaf;
    EDX ← MONITOR/MWAIT Leaf;
BREAK;
EAX = 6H:
    EAX ← Thermal and Power Management Leaf; (* See Table 2-8. *)
    EBX ← Thermal and Power Management Leaf;
    ECX ← Thermal and Power Management Leaf;
    EDX ← Thermal and Power Management Leaf;
BREAK;
EAX = 7H:
    EAX ← Structured Extended Feature Leaf; (* See Table 2-8. *);
    EBX ← Structured Extended Feature Leaf;
    ECX ← Structured Extended Feature Leaf;
    EDX ← Structured Extended Feature Leaf;
BREAK;
EAX = 8H:
    EAX ← Reserved = 0;

```



```

    EBX ← Reserved = 0;
    ECX ← Reserved = 0;
    EDX ← Reserved = 0;
BREAK;
EAX = 9H:
    EAX ← Direct Cache Access Information Leaf; (* See Table 2-8. *)
    EBX ← Direct Cache Access Information Leaf;
    ECX ← Direct Cache Access Information Leaf;
    EDX ← Direct Cache Access Information Leaf;
BREAK;
EAX = AH:
    EAX ← Architectural Performance Monitoring Leaf; (* See Table 2-8. *)
    EBX ← Architectural Performance Monitoring Leaf;
    ECX ← Architectural Performance Monitoring Leaf;
    EDX ← Architectural Performance Monitoring Leaf;
BREAK;
EAX = BH:
    EAX ← Extended Topology Enumeration Leaf; (* See Table 2-8. *)
    EBX ← Extended Topology Enumeration Leaf;
    ECX ← Extended Topology Enumeration Leaf;
    EDX ← Extended Topology Enumeration Leaf;
BREAK;
EAX = CH:
    EAX ← Reserved = 0;
    EBX ← Reserved = 0;
    ECX ← Reserved = 0;
    EDX ← Reserved = 0;
BREAK;
EAX = DH:
    EAX ← Processor Extended State Enumeration Leaf; (* See Table 2-8. *)
    EBX ← Processor Extended State Enumeration Leaf;
    ECX ← Processor Extended State Enumeration Leaf;
    EDX ← Processor Extended State Enumeration Leaf;
BREAK;
EAX = EH:
    EAX ← Reserved = 0;
    EBX ← Reserved = 0;
    ECX ← Reserved = 0;
    EDX ← Reserved = 0;
BREAK;
EAX = FH:
    EAX ← Platform Quality of Service Monitoring Enumeration Leaf; (* See Table 2-8. *)
    EBX ← Platform Quality of Service Monitoring Enumeration Leaf;
    ECX ← Platform Quality of Service Monitoring Enumeration Leaf;
    EDX ← Platform Quality of Service Monitoring Enumeration Leaf;
BREAK;
EAX = 10H:
    EAX ← Platform Quality of Service Enforcement Enumeration Leaf; (* See Table 2-8. *)
    EBX ← Platform Quality of Service Enforcement Enumeration Leaf;
    ECX ← Platform Quality of Service Enforcement Enumeration Leaf;
    EDX ← Platform Quality of Service Enforcement Enumeration Leaf;
BREAK;
EAX = 14H:
    EAX ← Intel Processor Trace Enumeration Leaf; (* See Table 2-8. *)

```

EBX ← Intel Processor Trace Enumeration Leaf;
 ECX ← Intel Processor Trace Enumeration Leaf;
 EDX ← Intel Processor Trace Enumeration Leaf;
 BREAK;
 EAX = 15H:
 EAX ← Time Stamp Counter and Core Crystal Clock Information Leaf; (* See Table 2-8. *)
 EBX ← Time Stamp Counter and Core Crystal Clock Information Leaf;
 ECX ← Time Stamp Counter and Core Crystal Clock Information Leaf;
 EDX ← Time Stamp Counter and Core Crystal Clock Information Leaf;
 BREAK;
 EAX = 16H:
 EAX ← Processor Frequency Information Enumeration Leaf; (* See Table 2-8. *)
 EBX ← Processor Frequency Information Enumeration Leaf;
 ECX ← Processor Frequency Information Enumeration Leaf;
 EDX ← Processor Frequency Information Enumeration Leaf;
 BREAK;
 EAX = 17H:
 EAX ← System-On-Chip Vendor Attribute Enumeration Leaf; (* See Table 2-8. *)
 EBX ← System-On-Chip Vendor Attribute Enumeration Leaf;
 ECX ← System-On-Chip Vendor Attribute Enumeration Leaf;
 EDX ← System-On-Chip Vendor Attribute Enumeration Leaf;
 BREAK;
 EAX = 18H:
 EAX ← Deterministic Address Translation Parameters Enumeration Leaf; (* See Table 2-8. *)
 EBX ← Deterministic Address Translation Parameters Enumeration Leaf;
 ECX ← Deterministic Address Translation Parameters Enumeration Leaf;
 EDX ← Deterministic Address Translation Parameters Enumeration Leaf;
 BREAK;
 EAX = 80000000H:
 EAX ← Highest extended function input value understood by CPUID;
 EBX ← Reserved;
 ECX ← Reserved;
 EDX ← Reserved;
 BREAK;
 EAX = 80000001H:
 EAX ← Reserved;
 EBX ← Reserved;
 ECX ← Extended Feature Bits (* See Table 2-8.*);
 EDX ← Extended Feature Bits (* See Table 2-8. *);
 BREAK;
 EAX = 80000002H:
 EAX ← Processor Brand String;
 EBX ← Processor Brand String, continued;
 ECX ← Processor Brand String, continued;
 EDX ← Processor Brand String, continued;
 BREAK;
 EAX = 80000003H:
 EAX ← Processor Brand String, continued;
 EBX ← Processor Brand String, continued;
 ECX ← Processor Brand String, continued;
 EDX ← Processor Brand String, continued;
 BREAK;
 EAX = 80000004H:
 EAX ← Processor Brand String, continued;

```

    EBX ← Processor Brand String, continued;
    ECX ← Processor Brand String, continued;
    EDX ← Processor Brand String, continued;
BREAK;
EAX = 80000005H:
    EAX ← Reserved = 0;
    EBX ← Reserved = 0;
    ECX ← Reserved = 0;
    EDX ← Reserved = 0;
BREAK;
EAX = 80000006H:
    EAX ← Reserved = 0;
    EBX ← Reserved = 0;
    ECX ← Cache information;
    EDX ← Reserved = 0;
BREAK;
EAX = 80000007H:
    EAX ← Reserved = 0;
    EBX ← Reserved = 0;
    ECX ← Reserved = 0;
    EDX ← Reserved = 0;
BREAK;
EAX = 80000008H:
    EAX ← Reserved = 0;
    EBX ← Reserved = 0;
    ECX ← Reserved = 0;
    EDX ← Reserved = 0;
BREAK;
DEFAULT: (* EAX = Value outside of recognized range for CPUID. *)
    (* If the highest basic information leaf data depend on ECX input value, ECX is honored. *)
    EAX ← Reserved; (* Information returned for highest basic information leaf. *)
    EBX ← Reserved; (* Information returned for highest basic information leaf. *)
    ECX ← Reserved; (* Information returned for highest basic information leaf. *)
    EDX ← Reserved; (* Information returned for highest basic information leaf. *)
BREAK;
ESAC;

```

Flags Affected

None.

Exceptions (All Operating Modes)

#UD If the LOCK prefix is used.

In earlier IA-32 processors that do not support the CPUID instruction, execution of the instruction results in an invalid opcode (#UD) exception being generated. §

CHAPTER 3

SYSTEM PROGRAMMING FOR INTEL® AVX-512

This chapter describes the operating system programming considerations for supporting the following extended processor states: 512-bit ZMM registers and opmask k-registers. These system programming requirements apply to AVX-512 Foundation instructions and other 512-bit instructions described in Chapter 6.

The basic requirements for an operating system using XSAVE/XRSTOR to manage processor extended states, e.g. YMM registers, can be found in Chapter 13 of *Intel 64 and IA-32 Architectures Software Developer's Manual, Volumes 3A*. This chapter covers additional requirements for OS to support ZMM and opmask register states.

3.1 AVX-512 STATE, EVEX PREFIX AND SUPPORTED OPERATING MODES

AVX-512 instructions are encoded using EVEX prefix. The EVEX encoding scheme can support 512-bit, 256-bit and 128-bit instructions that operate on opmask register, ZMM, YMM and XMM states.

For processors that support AVX-512 family of instructions, the extended processor states (ZMM and opmask registers) exist in all operating modes. However, the access to those states may vary in different modes. The processor's support for instruction extensions that employ EVEX prefix encoding is independent of the processor's support for using XSAVE/XRSTOR/XSAVEOPT to those states.

Instructions requiring EVEX prefix encoding generally are supported in 64-bit, 32-bit modes, and 16-bit protected mode. They are not supported in Real mode, Virtual-8086 mode or entering into SMM mode.

Note that bits MAX_VL-1:256 (511:256) of ZMM register state are maintained across transitions into and out of these modes. Because the XSAVE/XRSTOR/XSAVEOPT instruction can operate in all operating modes, it is possible that the processor's ZMM register state can be modified by software in any operating mode by executing XRSTOR. The ZMM registers can be updated by XRSTOR using the state information stored in the XSAVE/XRSTOR area residing in memory.

3.2 AVX-512 STATE MANAGEMENT

Operating systems must use the XSAVE/XRSTOR/XSAVEOPT instructions for ZMM and opmask state management. An OS must enable its ZMM and opmask state management to support AVX-512 Foundation instructions. Otherwise, an attempt to execute an instruction in AVX-512 Foundation instructions (including a scalar 128-bit SIMD instructions using EVEX encoding) will cause a #UD exception. An operating system, which enabled AVX-512 state to support AVX-512 Foundation instructions, is also sufficient to support the rest of AVX-512 family of instructions.

3.2.1 Detection of ZMM and Opmask State Support

Hardware support of the extended state components for executing AVX-512 Foundation instructions is queried through the main leaf of CPUID leaf function 0DH with index ECX = 0. Specifically, the return value in EDX:EAX of CPUID.(EAX=0DH, ECX=0) provides a 64-bit wide bit vector of hardware support of processor state components, beginning with bit 0 of EAX corresponding to x87 FPU state, CPUID.(EAX=0DH, ECX=0):EAX[1] corresponding to SSE state (XMM registers and MXCSR), CPUID.(EAX=0DH, ECX=0):EAX[2] corresponding to YMM states.

The ZMM and opmask states consist of three additional components in the XSAVE/XRSTOR state save area:

- The opmask register state component represents eight 64-bit opmask registers. Processor support for this component state is indicated by CPUID.(EAX=0DH, ECX=0):EAX[5].
- The ZMM_Hi256 component represents the high 256 bits of the low 16 ZMM registers, i.e. ZMM0..15[511:256]. Processor support for this component state is indicated by CPUID.(EAX=0DH, ECX=0):EAX[6].
- The Hi16_ZMM component represents the full 512 bits of the high 16 ZMM registers, i.e. ZMM16..31[511:0]. Processor support for this component state is indicated by CPUID.(EAX=0DH, ECX=0):EAX[7].

Each component state has a corresponding enable it in the XCR0 register. Operating system must use XSETBV to set these three enable bits to enable AVX-512 Foundation instructions to be decoded. The location of bit vector representing the AVX-512 states, matching the layout of the XCR0 register, is provided in the following figure.

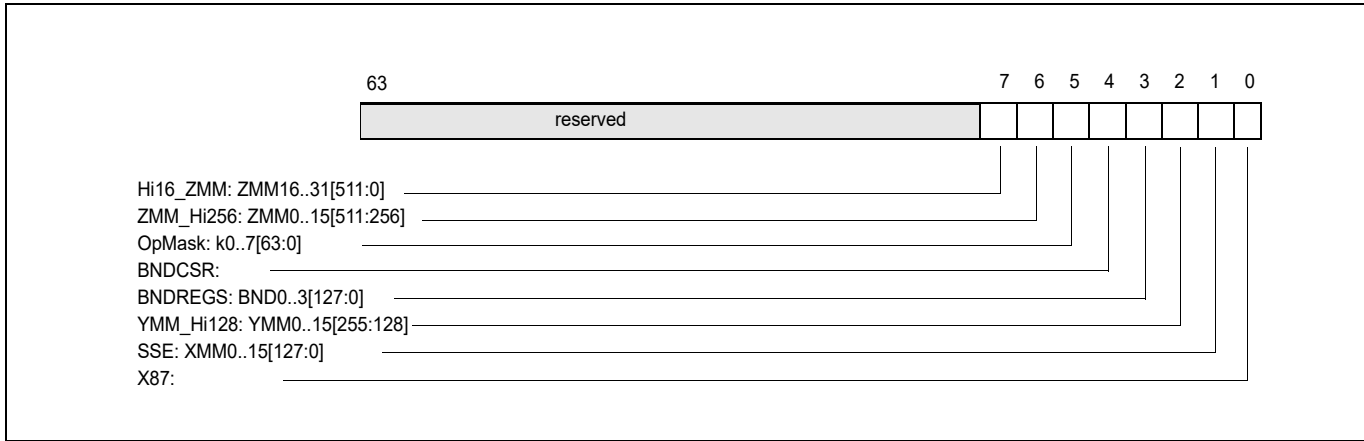


Figure 3-1. Bit Vector and XCR0 Layout of Extended Processor State Components

3.2.2 Enabling of ZMM and Opmask Register State

An OS can enable ZMM and opmask register state support with the following steps:

- Verify the processor supports XSAVE/XRSTOR/XSETBV/XGETBV instructions and the XCR0 register by checking CPUID.1.ECX.XSAVE[bit 26]=1.
- Verify the processor supports SSE, YMM, ZMM_Hi256, Hi16_ZMM, and opmask states (i.e. bits 2:1 and 7:5 of XCR0 are valid) by checking CPUID. (EAX=0DH, ECX=0):EAX[7:5].

The OS must determine the buffer size requirement for the XSAVE area that will be used by XSAVE/XRSTOR. Note that even though ZMM8-ZMM31 are not accessible in 32 bit mode, a 32 bit OS is still required to allocate the buffer for the entire ZMM state.

- Set CR4.OSXSAVE[bit 18]=1 to enable the use of XSETBV/XGETBV instructions to write/read the XCR0 register.
- Supply an appropriate mask via EDX:EAX to execute XSETBV to enable the processor state components that the OS wishes to manage using XSAVE/XRSTOR instruction.

To enable ZMM and opmask register state, system software must use a EDX:EAX mask of 111xx111b when executing XSETBV.

Table 3-1. XCR0 Processor State Components

Bit	Meaning
0 - x87	This bit 0 must be 1. An attempt to write 0 to this bit causes a #GP exception.
1 - SSE	If 1, the processor supports SSE state (MXCSR and XMM registers) management using XSAVE, XSAVEOPT, and XRSTOR. This bit must be set to '1' to enable AVX-512 Foundation instructions.
2 - YMM_Hi128	If 1, the processor supports YMM_hi128 state management (upper 128 bits of YMM0-15) using XSAVE, XSAVEOPT, and XRSTOR. This bit must be set to '1' to enable AVX-512 Foundation instructions.
3 - BNDREGS	If 1, the processor supports Intel Memory Protection Extensions (Intel MPX) bound register state management using XSAVE, XSAVEOPT, and XRSTOR.
4 - BNDCSR	If 1, the processor supports Intel MPX bound configuration and status management using XSAVE, XSAVEOPT, and XRSTOR.
5 - Opmask	If 1, the processor supports the opmask state management using XSAVE, XSAVEOPT, and XRSTOR. This bit must be set to '1' to enable AVX-512 Foundation instructions.

Table 3-1. XCR0 Processor State Components

Bit	Meaning
6 - ZMM_Hi256	If 1, the processor supports ZMM_Hi256 state (the upper 256 bits of the low 16 ZMM registers) management using XSAVE, XSAVEOPT, and XRSTOR. This bit must be set to '1' to enable AVX-512 Foundation instructions.
7 - Hi16_ZMM	If 1, the processor supports Hi16_ZMM state (the full 512 bits of the high16 ZMM registers) management using XSAVE, XSAVEOPT, and XRSTOR. This bit must be set to '1' to enable AVX-512 Foundation instructions.

3.2.3 Enabling of SIMD Floating-Exception Support

AVX-512 Foundation instructions may generate SIMD floating-point exceptions. An OS must enable SIMD floating-point exception support by setting CR4.OSXMMEXCPT[bit 10]=1.

The effect of CR4 setting that affects AVX-512 Foundation instructions is the same as for AVX and FMA enabling as listed in Table 3-2

Table 3-2. CR4 Bits for AVX-512 Foundation Instructions Technology Support

Bit	Meaning
CR4.OSXSAVE[bit 18]	If set, the OS supports use of XSETBV/XGETBV instruction to access the XCR0 register, XSAVE/XRSTOR to manage processor extended states. Must be set to '1' to enable AVX-512 Foundation, AVX2, FMA, and AVX instructions.
CR4.OSXMMEXCPT[bit 10]	Must be set to 1 to enable SIMD floating-point exceptions. This applies to SIMD floating-point instructions across AVX-512 Foundation, AVX and FMA, and legacy 128-bit SIMD floating-point instructions operating on XMM registers.
CR4.OSFXSR[bit 9]	Must be set to 1 to enable legacy 128-bit SIMD instructions operating on XMM state. Not needed to enable AVX-512 Foundation, AVX2, FMA, and AVX instructions.

3.2.4 The Layout of XSAVE State Save Area

The OS must determine the buffer size requirement by querying CPUID with EAX=0DH, ECX=0. If the OS wishes to enable all processor extended state components in the XCR0, it can allocate the buffer size according to CPUID.(EAX=0DH, ECX=0):ECX.

After the memory buffer for XSAVE is allocated, the entire buffer must be cleared prior to executing XSAVE.

The XSAVE area layout currently defined in Intel Architecture is listed in Table 3-3. The register fields of the first 512 byte of the XSAVE area are identical to those of the FXSAVE/FXRSTOR area.

The layout of the XSAVE Area for additional processor components (512-bit ZMM register, 32 ZMM registers, opmask registers) are to be determined later.

Table 3-3. Layout of XSAVE Area For Processor Supporting YMM State

Save Areas	Offset (Byte)	Size (Bytes)
FPU/SSE SaveArea	0	512
Header	512	64
Ext_Save_Area_2 (YMM_Hi128)	CPUID.(EAX=0DH, ECX=2):EBX	CPUID.(EAX=0DH, ECX=2):EAX
Ext_Save_Area_3 (BNDREGS)	CPUID.(EAX=0DH, ECX=3):EBX	CPUID.(EAX=0DH, ECX=3):EAX
Ext_Save_Area_4 (BNDCSR)	CPUID.(EAX=0DH, ECX=4):EBX	CPUID.(EAX=0DH, ECX=4):EAX
Ext_Save_Area_5 (OPMASK)	CPUID.(EAX=0DH, ECX=5):EBX	CPUID.(EAX=0DH, ECX=5):EAX
Ext_Save_Area_6 (ZMM_Hi256)	CPUID.(EAX=0DH, ECX=6):EBX	CPUID.(EAX=0DH, ECX=6):EAX
Ext_Save_Area_7 (Hi16_ZMM)	CPUID.(EAX=0DH, ECX=7):EBX	CPUID.(EAX=0DH, ECX=7):EAX

The format of the header is as follows (see Table 3-4):

Table 3-4. XSAVE Header Format

15:8	7:0	Byte Offset from Header	Byte Offset from XSAVE Area
Reserved (Must be zero)	XSTATE_BV	0	512
Reserved	Reserved (Must be zero)	16	528
Reserved	Reserved	32	544
Reserved	Reserved	48	560

The layout of the Ext_Save_Area[YMM_Hi128] contains 16 of the upper 128-bits of the YMM registers, it is shown in Table 3-5.

Table 3-5. XSAVE Save Area Layout for YMM_Hi128 State (Ext_Save_Area_2)

31 16	15 0	Byte Offset from YMM_Hi128_Save_Area	Byte Offset from XSAVE Area
YMM1[255:128]	YMM0[255:128]	0	576
YMM3[255:128]	YMM2[255:128]	32	608
YMM5[255:128]	YMM4[255:128]	64	640
YMM7[255:128]	YMM6[255:128]	96	672
YMM9[255:128]	YMM8[255:128]	128	704
YMM11[255:128]	YMM10[255:128]	160	736
YMM13[255:128]	YMM12[255:128]	192	768
YMM15[255:128]	YMM14[255:128]	224	800

The layout of the Ext_SAVE_Area_3[BNDREGS] contains bounds register state of the Intel Memory Protection Extensions (Intel MPX).

The layout of the Ext_SAVE_Area_4[BNDCSR] contains the processor state of bounds configuration and status of Intel MPX.

The layout of the Ext_SAVE_Area_5[Opmask] contains 8 64-bit mask register as shown in Table 3-6.

Table 3-6. XSAVE Save Area Layout for Opmask Registers

15 8	7 0	Byte Offset from OPMASK_Save_Area	Byte Offset from XSAVE Area
K1[63:0]	K0[63:0]	0	1088
K3[63:0]	K2[63:0]	16	1104
K5[63:0]	K4[63:0]	32	1120
K7[63:0]	K6[63:0]	48	1136

The layout of the Ext_SAVE_Area_6[ZMM_Hi256] is shown below in Table 3-7.

Table 3-7. XSAVE Save Area Layout for ZMM State of the High 256 Bits of ZMM0-ZMM15 Registers

63 32	31 0	Byte Offset from ZMM_Hi256_Save_Area	Byte Offset from XSAVE Area
ZMM1[511:256]	ZMM0[511:256]	0	1152
ZMM3[511:256]	ZMM2[511:256]	64	1216
ZMM5[511:256]	ZMM4[511:256]	128	1280
ZMM7[511:256]	ZMM6[511:256]	192	1344
ZMM9[511:256]	ZMM8[511:256]	256	1408
ZMM11[511:256]	ZMM10[511:256]	320	1472
ZMM13[511:256]	ZMM12[511:256]	384	1536
ZMM15[511:256]	ZMM14[511:256]	448	1600

The layout of the Ext_SAVE_Area_7[Hi16_ZMM] corresponding to the upper new 16 ZMM registers is shown below in Table 3-8.

Table 3-8. XSAVE Save Area Layout for ZMM State of ZMM16-ZMM31 Registers

127 64	63 0	Byte Offset from Hi16_ZMM_Save_Area	Byte Offset from XSAVE Area
ZMM17[511:0]	ZMM16[511:0]	0	1664
ZMM19[511:0]	ZMM18[511:0]	128	1792
ZMM21[511:0]	ZMM20[511:0]	256	1920
ZMM23[511:0]	ZMM22[511:0]	384	2048
ZMM25[511:0]	ZMM24[511:0]	512	2176
ZMM27[511:0]	ZMM26[511:0]	640	2304
ZMM29[511:0]	ZMM28[511:0]	768	2432
ZMM31[511:0]	ZMM30[511:0]	896	2560

3.2.5 XSAVE/XRSTOR Interaction with YMM State and MXCSR

The processor's actions as a result of executing XRSTOR, on the MXCSR, XMM and YMM registers, are listed in Table 3-9. The XMM registers may be initialized by the processor (See XRSTOR operation in *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B*). When the MXCSR register is updated from memory, reserved bit checking is enforced. XSAVE / XRSTOR will save / restore the MXCSR only if the AVX or SSE bits are set in the EDX:EAX mask.

Table 3-9. XRSTOR Action on MXCSR, XMM Registers, YMM Registers

EDX:EAX		XSTATE_BV		MXCSR	YMM_Hi128 Registers	XMM Registers
Bit 2	Bit 1	Bit 2	Bit 1			
0	0	X	X	None	None	None
0	1	X	0	Load/Check	None	Init by processor
0	1	X	1	Load/Check	None	Load
1	0	0	X	Load/Check	Init by processor	None
1	0	1	X	Load/Check	Load	None
1	1	0	0	Load/Check	Init by processor	Init by processor
1	1	0	1	Load/Check	Init by processor	Load
1	1	1	0	Load/Check	Load	Init by processor
1	1	1	1	Load/Check	Load	Load

The action of XSAVE for managing YMM and MXCSR is listed in Table 3-10.

Table 3-10. XSAVE Action on MXCSR, XMM, YMM Register

EDX:EAX		XCRO_MASK		MXCSR	YMM_H Registers	XMM Registers
Bit 2	Bit 1	Bit 2	Bit 1			
0	0	X	X	None	None	None
0	1	X	1	Store	None	Store
0	1	X	0	None	None	None
1	0	0	X	None	None	None
1	0	1	1	Store	Store	None
1	1	0	0	None	None	None
1	1	0	1	Store	None	Store
1	1	1	1	Store	Store	Store

3.2.6 XSAVE/XRSTOR/XSAVEOPT and Managing ZMM and Opmask States

The requirements for managing ZMM_Hi256, Hi16_ZMM and Opmask registers using XSAVE/XRSTOR/XSAVEOPT are simpler than those listed in Section 3.2.5. Because each of the three components (ZMM_Hi256, Hi16_ZMM and Opmask registers) can be managed independently of one another by XSAVE/XRSTOR/XSAVEOPT according to the corresponding bits in the bit vectors: EDX:EAX, XSAVE_BV, XCRO_MASK, independent of MXCSR:

- For using XSAVE with Opmask/ZMM_Hi256/Hi16_ZMM, XSAVE/XSAVEOPT will save the component to memory and mark the corresponding bits in the XSTATE_BV of the XSAVE header, if that component is specified in EDX:EAX as input to XSAVE/XSAVEOPT.
- XRSTOR will restore the Opmask/ZMM_Hi256/Hi16_ZMM components by checking the corresponding bits in both the input bit vector in EDX:EAX of XRSTOR and in XSTATE_BV of the header area in the following ways:
 - If the corresponding bit in EDX:EAX is set and XSTATE_BV is INIT, that component will be initialized,
 - If the corresponding bit in EDX:EAX is set and XSTATE_BV is set, that component will be restored from memory,
 - If the corresponding bit in EDX:EAX is not set, that component will remain unchanged.
- To enable AVX-512 Foundation instructions, all three components (Opmask/ZMM_Hi256/Hi16_ZMM) in XCRO must be set.

The processor supplied INIT values for each processor state component used by XRSTOR is listed in Table 3-11.

Table 3-11. Processor Supplied Init Values XRSTOR May Use

Processor State Component	Processor Supplied Register Values
x87 FPU State	FCW \leftarrow 037FH; FTW \leftarrow 0FFFFH; FSW \leftarrow 0H; FPU CS \leftarrow 0H; FPU DS \leftarrow 0H; FPU IP \leftarrow 0H; FPU DP \leftarrow 0; ST0-ST7 \leftarrow 0;
SSE State ¹	If 64-bit Mode: XMM0-XMM15 \leftarrow 0H; Else XMM0-XMM7 \leftarrow 0H
YMM_Hi128 State ¹	If 64-bit Mode: YMM0_H-YMM15_H \leftarrow 0H; Else YMM0_H-YMM7_H \leftarrow 0H
OPMASK State ¹	If 64-bit Mode: K0-K7 \leftarrow 0H;
ZMM_Hi256 State ¹	If 64-bit Mode: ZMM0_H-ZMM15_H \leftarrow 0H; Else ZMM0_H-ZMM7_H \leftarrow 0H
Hi16_ZMM State ¹	If 64-bit Mode: ZMM16-ZMM31 \leftarrow 0H;

NOTES:

1. MXCSR state is not updated by processor supplied values. MXCSR state can only be updated by XRSTOR from state information stored in XSAVE/XRSTOR area.

3.3 RESET BEHAVIOR

At processor reset

- YMM0-15 bits[255:0] are set to zero.
- ZMM0-15 bits [511:256] are set to zero.
- ZMM16-31 are set to zero.
- Opmask register K0-7 are set to 0x0H.
- **XCRO**[2:1] is set to zero, **XCRO**[0] is set to 1.
- **XCRO**[7:6] and is set to zero, **XCRO**[Opmask] is set to 0.
- CR4.OSXSAVE[bit 18] (and its mirror CPUID.1.ECX.OSXSAVE[bit 27]) is set to 0.

3.4 EMULATION

Setting the CR0.EM bit to 1 provides a technique to emulate Legacy SSE floating-point instruction sets in software. This technique is not supported with AVX instructions, nor FMA instructions.

If an operating system wishes to emulate AVX instructions, set XCRO[2:1] to zero. This will cause AVX instructions to #UD. Emulation of FMA by operating system can be done similarly as with emulating AVX instructions.

3.5 WRITING FLOATING-POINT EXCEPTION HANDLERS

AVX-512, AVX and FMA floating-point exceptions are handled in an entirely analogous way to Legacy SSE floating-point exceptions. To handle unmasked SIMD floating-point exceptions, the operating system or executive must provide an exception handler. The section titled “SSE and SSE2 SIMD Floating-Point Exceptions” in Chapter 11, “Programming with Streaming SIMD Extensions 2 (SSE2),” of the *IA-32 Intel® Architecture Software Developer’s Manual, Volume 1*, describes the SIMD floating-point exception classes and gives suggestions for writing an exception handler to handle them.

To indicate that the operating system provides a handler for SIMD floating-point exceptions (#XM), the CR4.OSXM-MEXCPT flag (bit 10) must be set.

CHAPTER 4

INTEL® AVX-512 INSTRUCTION ENCODING

4.1 OVERVIEW SECTION

This chapter describes the details of the Intel® AVX-512 instruction encoding system. The AVX-512 Foundation instructions described in Chapter 5 use a new prefix (called EVEX). Opmask instructions described in Chapter 6 are encoded using the VEX prefix. The EVEX prefix has some parts resembling the instruction encoding scheme using the VEX prefix, and many other capabilities not available with the VEX prefix. The EVEX encoding architecture also applies to other 512-bit instructions described in Chapter 6.

The significant feature differences between EVEX and VEX are summarized below.

- EVEX is a 4-Byte prefix (the first byte must be 62H); VEX is either a 2-Byte (C5H is the first byte) or 3-Byte (C4H is the first byte) prefix.
- EVEX prefix can encode 32 vector registers (XMM/YMM/ZMM) in 64-bit mode.
- EVEX prefix can encode an opmask register for conditional processing or selection control in EVEX-encoded vector instructions; opmask instructions, whose source/destination operands are opmask registers and treat the content of an opmask register as a single value, are encoded using the VEX prefix.
- EVEX memory addressing with disp8 form uses a compressed disp8 encoding scheme to improve encoding density of the instruction byte stream.
- EVEX prefix can encode functionality that are specific to instruction classes (e.g. packed instruction with “load+op” semantic can support embedded broadcast functionality, floating-point instruction with rounding semantic can support static rounding functionality, floating-point instruction with non-rounding arithmetic semantic can support “suppress all exceptions” functionality).

4.2 INSTRUCTION FORMAT AND EVEX

The placement of the EVEX prefix in an IA instruction is represented in Figure 4-1:




Figure 4-1. AVX-512 Instruction Format and the EVEX Prefix

The EVEX prefix is a 4-byte prefix, with the first two bytes derived from unused encoding form of the 32-bit-mode-only BOUND instruction. The layout of the EVEX prefix is shown in Figure 4-2. The first byte must be 62H, followed by three payload bytes, denoted as P0, P1, and P2 individually or collectively as P[23:0] (see Figure 4-2).

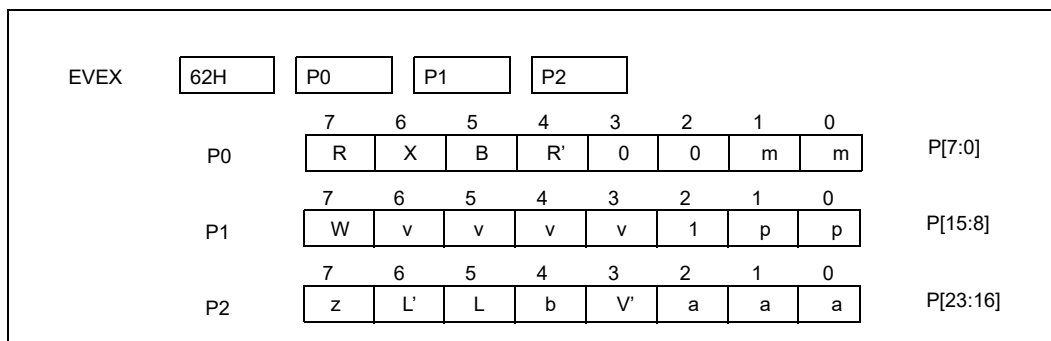


Figure 4-2. Bit Field Layout of the EVEX Prefix

Table 4-1. EVEX Prefix Bit Field Functional Grouping

Notation	Bit field Group	Position	Comment
--	Reserved	P[3 : 2]	Must be 0
--	Fixed Value	P[10]	Must be 1
EVEX.mm	Compressed legacy escape	P[1: 0]	Identical to low two bits of VEX.mmmmm
EVEX.pp	Compressed legacy prefix	P[9 : 8]	Identical to VEX.pp
EVEX.RXB	Next-8 register specifier modifier	P[7 : 5]	Combine with ModR/M.reg, ModR/M.rm (base, index/vidx)
EVEXR'	High-16 register specifier modifier	P[4]	Combine with EVEX.R and ModR/M.reg
EVEXX	High-16 register specifier modifier	P[6]	Combine with EVEX.B and ModR/M.rm, when SIB/VSIB absent
EVEX.vvvv	NDS register specifier	P[14 : 11]	Same as VEX.vvvv
EVEXV'	High-16 NDS/VIDX register specifier	P[19]	Combine with EVEX.vvvv or when VSIB present
EVEX.aaa	Embedded opmask register specifier	P[18 : 16]	
EVEX.W	Osize promotion/Opcode extension	P[15]	
EVEX.z	Zeroing/Merging	P[23]	
EVEX.b	Broadcast/RC/SAE Context	P[20]	
EVEX.L'L	Vector length/RC	P[22 : 21]	

The bit fields in P[23:0] are divided into the following functional groups (Table 4-1 provides a tabular summary):

- Reserved bits: P[3:2] must be 0, otherwise #UD.
- Fixed-value bit: P[10] must be 1, otherwise #UD.
- Compressed legacy prefix/escape bytes: P[1:0] is identical to the lowest 2 bits of VEX.mmmmm; P[9:8] is identical to VEX.pp.
- Operand specifier modifier bits for vector register, general purpose register, memory addressing: P[7:5] allows access to the next set of 8 registers beyond the low 8 registers when combined with ModR/M register specifiers.
- Operand specifier modifier bit for vector register: P[4] (or EVEX.R') allows access to the high 16 vector register set when combined with P[7] and ModR/M.reg specifier; P[6] can also provide access to a high 16 vector register when SIB or VSIB addressing are not needed.
- Non-destructive source /vector index operand specifier: P[19] and P[14:11] encode the second source vector register operand in a non-destructive source syntax, vector index register operand can access an upper 16 vector register using P[19].
- Op-mask register specifiers: P[18:16] encodes op-mask register set k0-k7 in instructions operating on vector registers.

- EVEX.W: P[15] is similar to VEX.W which serves either as opcode extension bit or operand size promotion to 64-bit in 64-bit mode.
- Vector destination merging/zeroing: P[23] encodes the destination result behavior which either zeroes the masked elements or leave masked element unchanged.
- Broadcast/Static-rounding/SAE context bit: P[20] encodes multiple functionality, which differs across different classes of instructions and can affect the meaning of the remaining field (EVEX.L'L). The functionality for the following instruction classes are:
 - Broadcasting a single element across the destination vector register: this applies to the instruction class with Load+Op semantic where one of the source operand is from memory.
 - Redirect L'L field (P[22:21]) as static rounding control for floating-point instructions with rounding semantic. Static rounding control overrides MXCSR.RC field and implies "Suppress all exceptions" (SAE).
 - Enable SAE for floating -point instructions with arithmetic semantic that is not rounding.
 - For instruction classes outside of the afore-mentioned three classes, setting EVEX.b will cause #UD.
- Vector length/rounding control specifier: P[22:21] can server one of three functionality:
 - vector length information for packed vector instructions,
 - ignored for instructions operating on vector register content as a single data element,
 - rounding control for floating-point instructions that have a rounding semantic and whose source and destination operands are all vector registers.

4.3 REGISTER SPECIFIER ENCODING AND EVEX

EVEX-encoded instruction can access 8 opmask registers, 16 general-purpose registers and 32 vector registers in 64-bit mode (8 general-purpose registers and 8 vector registers in non-64-bit modes). EVEX-encoding can support instruction syntax that access up to 4 instruction operands. Normal memory addressing modes and VSIB memory addressing are supported with EVEX prefix encoding. The mapping of register operands used by various instruction syntax and memory addressing in 64-bit mode are shown in Table 4-2. Opmask register encoding is described in Section 4.3.1.

Table 4-2. 32-Register Support in 64-bit Mode Using EVEX with Embedded REX Bits

	4 ¹	3	[2:0]	Reg. Type	Common Usages
REG	EVEX.R'	REX.R	modrm.reg	GPR, Vector	Destination or Source
NDS/NDD	EVEX.V'	EVEX.vvvv		GPR, Vector	2ndSource or Destination
RM	EVEX.X	EVEX.B	modrm.r/m	GPR, Vector	1st Source or Destination
BASE	0	EVEX.B	modrm.r/m	GPR	memory addressing
INDEX	0	EVEX.X	sib.index	GPR	memory addressing
VIDX	EVEX.V'	EVEX.X	sib.index	Vector	VSIB memory addressing
IS4	Imm8[3]	Imm8[7:4]		Vector	3rd Source

NOTES:

1. Not applicable for accessing general purpose registers.

The mapping of register operands used by various instruction syntax and memory addressing in 32-bit modes are shown in Table 4-3.

Table 4-3. EVEX Encoding Register Specifiers in 32-bit Mode

	[2:0]	Reg. Type	Common Usages
REG	modrm.reg	GPR, Vector	Dest or Source
NDS/NDD	EVEX.vvv	GPR, Vector	2ndSource or Dest
RM	modrm.r/m	GPR, Vector	1st Source or Dest
BASE	modrm.r/m	GPR	memory addressing
INDEX	sib.index	GPR	memory addressing
VIDX	sib.index	Vector	VSIB memory addressing
IS4	Imm8[7:5]	Vector	3rd Source

4.3.1 Opmask Register Encoding

There are eight opmask registers, k0-k7. Opmask register encoding falls into two categories:

- Opmask registers that are the source or destination operands of an instruction treating the content of opmask register as a scalar value, are encoded using the VEX prefix scheme. It can support up to three operands using standard modR/M byte's reg field and rm field and VEX.vvvv. Such a scalar opmask instruction does not support conditional update of the destination operand.
- An opmask register providing conditional processing and/or conditional update of the destination register of a vector instruction is encoded using EVEX.aaa field (see Section 4.4).
- An opmask register serving as the destination or source operand of a vector instruction is encoded using standard modR/M byte's reg field and rm fields.

Table 4-4. Opmask Register Specifier Encoding

	[2:0]	Register Access	Common Usages
REG	modrm.reg	k0-k7	Source
NDS	VEX.vvv	k0-k7	2ndSource
RM	modrm.r/m	k0-7	1st Source
{k1}	EVEX.aaa	k0 ¹ -k7	Opmask

NOTES:

1. instructions that overwrite the conditional mask in opmask do not permit using k0 as the embedded mask.

4.4 MASKING SUPPORT IN EVEX

EVEX can encode an opmask register to conditionally control per-element computational operation and updating of result of an instruction to the destination operand. The predicate operand is known as the opmask register. The EVEX.aaa field, P[18:16] of the EVEX prefix, is used to encode one out of a set of eight 64-bit architectural registers. Note that from this set of 8 architectural registers, only k1 through k7 can be addressed as predicate operands. k0 can be used as a regular source or destination but cannot be encoded as a predicate operand.

AVX-512 instructions support two types of masking with EVEX.z bit (P[23]) controlling the type of masking:

- Merging-masking, which is the default type of masking for EVEX-encoded vector instructions, preserves the old value of each element of the destination where the corresponding mask bit has a 0. It corresponds to the case of EVEX.z = 0.

- Zeroing-masking, is enabled by having the EVEX.z bit set to 1. In this case, an element of the destination is set to 0 when the corresponding mask bit has a 0 value.

AVX-512 Foundation instructions can be divided in three different groups:

- Instructions which support “zeroing-masking”.
 - Also allow merging-masking.
- Instructions which require aaa = 000.
 - Do not allow any form of masking.
- Instructions which allow merging-masking but do not allow zeroing-masking
 - Require EVEX.z to be set to 0
 - This group is mostly composed of instructions that write to memory.
- Instructions which require aaa <> 000 do not allow EVEX.z to be set to 1.
 - Allow merging-masking and do not allow zeroing-masking, e.g., gather instructions.

4.5 COMPRESSED DISPLACEMENT (DISP8*N) SUPPORT IN EVEX

For memory addressing using disp8 form, EVEX-encoded instructions always use a compressed displacement scheme by multiplying disp8 in conjunction with a scaling factor N that is determined based on the vector length, the value of EVEX.b bit (embedded broadcast) and the input element size of the instruction. In general, the factor N corresponds to the number of bytes characterizing the internal memory operation of the input operand (e.g., 64 when the accessing a full 512-bit memory vector). The scale factor N is listed in Table 4-5 and Table 4-6 below, where EVEX encoded instructions are classified using the **tupletype** attribute. The scale factor N of each tupletype is listed based on the vector length (VL) and other factors affecting it.

Table 4-5 covers EVEX-encoded instructions which has a load semantic in conjunction with additional computational or data element movement operation, operating either on the full vector or half vector (due to conversion of numerical precision from a wider format to narrower format). EVEX.b is supported for such instructions for data element sizes which are either dword or qword (see Section 4.7).

EVEX-encoded instruction that are pure load/store, and “Load+op” instruction semantic that operate on data element size less than dword do not support broadcasting using EVEX.b. These are listed in Table 4-6. Table 4-6 also includes many broadcast instructions which perform broadcast using a subset of data elements without using EVEX.b. These instructions and a few data element size conversion instruction are covered in Table 4-6. Instruction classified in Table 4-6 do not use EVEX.b and EVEX.b must be 0, otherwise #UD will occur.

The tupletype abbreviation will be referenced in the instruction operand encoding table in the reference page of each instruction, providing the cross reference for the scaling factor N to encoding memory addressing operand.

Note that the disp8*N rules still apply when using 16b addressing.

Table 4-5. Compressed Displacement (DISP8*N) Affected by Embedded Broadcast

TupleType	EVEX.b	InputSize	EVEX.W	Broadcast	N (VL=128)	N (VL=256)	N (VL= 512)	Comment
Full Vector (FV)	0	32bit	0	none	16	32	64	Load+Op (Full Vector Dword/Qword)
	1	32bit	0	{1tox}	4	4	4	
	0	64bit	1	none	16	32	64	
	1	64bit	1	{1tox}	8	8	8	
Half Vector (HV)	0	32bit	0	none	8	16	32	Load+Op (Half Vector)
	1	32bit	0	{1tox}	4	4	4	

Table 4-6. EVEX DISP8*N For Instructions Not Affected by Embedded Broadcast

TupleType	InputSize	EVEX.W	N (VL= 128)	N (VL= 256)	N (VL= 512)	Comment
Full Vector Mem (FVM)	N/A	N/A	16	32	64	Load/store or subDword full vector
Tuple1 Scalar (T1S)	8bit	N/A	1	1	1	1 Tuple less than Full Vector
	16bit	N/A	2	2	2	
	32bit	0	4	4	4	
	64bit	1	8	8	8	
Tuple1 Fixed (T1F)	32bit	N/A	4	4	4	1 Tuple memsize not affected by EVEX.W
	64bit	N/A	8	8	8	
Tuple2 (T2)	32bit	0	8	8	8	Broadcast (2 elements)
	64bit	1	NA	16	16	
Tuple4 (T4)	32bit	0	NA	16	16	Broadcast (4 elements)
	64bit	1	NA	NA	32	
Tuple8 (T8)	32bit	0	NA	NA	32	Broadcast (8 elements)
Half Mem (HVM)	N/A	N/A	8	16	32	SubQword Conversion
QuarterMem (QVM)	N/A	N/A	4	8	16	SubDword Conversion
OctMem (OVM)	N/A	N/A	2	4	8	SubWord Conversion
Mem128 (M128)	N/A	N/A	16	16	16	Shift count from memory
MOVDDUP (DUP)	N/A	N/A	8	32	64	VMOVDDUP

4.6 EVEX ENCODING OF BROADCAST/ROUNDING/SAE SUPPORT

EVEX.b can provide three types of encoding context, depending on the instruction classes:

- Embedded broadcasting of one data element from a source memory operand to the destination for vector instructions with “load+op” semantic.
- Static rounding control overriding MXCSR.RC for floating-point instructions with rounding semantic.
- “Suppress All exceptions” (SAE) overriding MXCSR mask control for floating-point arithmetic instructions that do not have rounding semantic.

4.6.1 Embedded Broadcast Support in EVEX

EVEX encodes an embedded broadcast functionality that is supported on many vector instructions with 32-bit (double word or single-precision floating-point) and 64-bit data elements, and when the source operand is from memory. EVEX.b (P[20]) bit is used to enable broadcast on load-op instructions. When enabled, only one element is loaded from memory and broadcasted to all other elements instead of loading the full memory size.

The following instruction classes do not support embedded broadcasting:

- Instructions with only one scalar result is written to the vector destination.
- Instructions with explicit broadcast functionality provided by its opcode.
- Instruction semantic is a pure load or a pure store operation.

4.6.2 Static Rounding Support in EVEX

Static rounding control embedded in the EVEX encoding system applies only to register-to-register flavor of floating-point instructions with rounding semantic at two distinct vector lengths: (i) scalar, (ii) 512-bit. In both

cases, the field EVEX.L'L expresses rounding mode control overriding MXCSR.RC if EVEX.b is set. When EVEX.b is set, “suppress all exceptions” is implied. The processor behaves as if all MXCSR masking controls are set, and none of the MXCSR flags will be updated.

4.6.3 SAE Support in EVEX

The EVEX encoding system allows arithmetic floating-point instructions without rounding semantic to be encoded with the SAE attribute. This capability applies to scalar and 512-bit vector lengths, register-to-register only, by setting EVEX.b. When EVEX.b is set, “suppress all exceptions” is implied. The processor behaves as if all MXCSR masking controls are set, and none of the MXCSR flags will be updated.

4.6.4 Vector Length Orthogonality

The architecture of EVEX encoding scheme can support SIMD instructions operating at multiple vector lengths. Many AVX-512 Foundation instructions operate at 512-bit vector length. The vector length of EVEX encoded vector instructions are generally determined using the L'L field in EVEX prefix, except for 512-bit floating-point, reg-reg instructions with rounding semantic. The table below shows the vector length corresponding to various values of the L'L bits. When EVEX is used to encode scalar instructions, L'L is generally ignored.

When EVEX.b bit is set for a register-register instructions with floating-point rounding semantic, the same two bits P2[6:5] specifies rounding mode for the instruction, with implied SAE behavior. The mapping of different instruction classes relative to the embedded broadcast/rounding/SAE control and the EVEX.L'L fields are summarized in Table 4-7.

Table 4-7. EVEX Embedded Broadcast/Rounding/SAE and Vector Length on Vector Instructions

Position	P2[4]	P2[6:5]	P2[6:5]
Broadcast/Rounding/SAE Context	EVEX.b	EVEX.L'L	EVEX.RC
Reg-reg, FP Instructions w/ rounding semantic	Enable static rounding control (SAE implied)	Vector length Implied (512 bit or scalar)	00b: SAE + RNE 01b: SAE + RD 10b: SAE + RU 11b: SAE + RZ
FP Instructions w/o rounding semantic, can cause #XF	SAE control	00b: 128-bit 01b: 256-bit 10b: 512-bit 11b: Reserved (#UD)	NA
Load+op Instructions w/ memory source	Broadcast Control		NA
Other Instructions (Explicit Load/Store/Broadcast/Gather/Scatter)	Must be 0 (otherwise #UD)		NA

4.7 #UD EQUATIONS FOR EVEX

Instructions encoded using EVEX can face three types of UD conditions: state dependent, opcode independent and opcode dependent.

4.7.1 State Dependent #UD

In general, attempts to execute an instruction, which required OS support for incremental extended state component, will #UD if required state components were not enabled by OS. Table 4-8 lists instruction categories with respect to required processor state components. Attempts to execute a given category of instructions while enabled states were less than the required bit vector in XCRO shown in Table 4-8 will cause #UD.

Table 4-8. OS XSAVE Enabling Requirements of Instruction Categories

Instruction Categories	Vector Register State Access	Required XCR0 Bit Vector [7:0]
Legacy SIMD prefix encoded Instructions (e.g SSE)	XMM	xxxxxx11b
VEX-encoded instructions operating on YMM	YMM	xxxxx111b
EVEX-encoded 128-bit instructions	ZMM	111xx111b
EVEX-encoded 256-bit instructions	ZMM	111xx111b
EVEX-encoded 512-bit instructions	ZMM	111xx111b
VEX-encoded instructions operating on opmask	k-reg	xx1xxx11b

4.7.2 Opcode Independent #UD

A number of bit fields in EVEX encoded instruction must obey mode-specific but opcode-independent patterns listed in Table 4-9:

Table 4-9. Opcode Independent, State Dependent EVEX Bit Fields

Position	Notation	64-bit #UD	Non-64-bit #UD
P[3 : 2]	--	if > 0	if > 0
P[10]	--	if 0	if 0
P[1 : 0]	EVEX.mm	if 00b	if 00b
P[7 : 6]	EVEX.RX	None (valid)	None (BOUND if EVEX.RX != 11b)

4.7.3 Opcode Dependent #UD

This section describes legal values for the rest of the EVEX bit fields. Table 4-10 lists the #UD conditions of EVEX prefix bit fields which encodes or modifies register operands.

Table 4-10. #UD Conditions of Operand-Encoding EVEX Prefix Bit Fields

Notation	Position	Operand Encoding	64-bit #UD	Non-64-bit #UD
EVEX.R	P[7]	ModRM.reg encodes k-reg	if EVEX.R = 0	None (BOUND if EVEX.RX != 11b)
		ModRM.reg is opcode extension	None (ignored)	
		ModRM.reg encodes all other registers	None (valid)	
EVEX.X	P[6]	ModRM.r/m encodes ZMM/YMM/XMM	None (valid)	
		ModRM.r/m encodes k-reg or GPR	None (ignored)	
		ModRM.r/m without SIB/VSIB	None (ignored)	
		ModRM.r/m with SIB/VSIB	None (valid)	
EVEX.B	P[5]	ModRM.r/m encodes k-reg	None (ignored)	None (ignored)
		ModRM.r/m encodes other registers	None (valid)	
		ModRM.r/m base present	None (valid)	
		ModRM.r/m base not present	None (ignored)	
EVEXR'	P[4]	ModRM.reg encodes k-reg or GPR	if 0	None (ignored)
		ModRM.reg is opcode extension	None (ignored)	
		ModRM.reg encodes ZMM/YMM/XMM	None (valid)	

Table 4-10. #UD Conditions of Operand-Encoding EVEX Prefix Bit Fields (Continued)

EVEX.vvvv	P[14 : 11]	vvvv encodes ZMM/YMM/XMM	None (valid)	None (valid) P[14] ignored
		otherwise	if != 1111b	if != 1111b
EVEX.v'	P[19]	encodes ZMM/YMM/XMM	None (valid)	if 0
		otherwise	if 0	if 0

Table 4-11 lists the #UD conditions of instruction encoding of opmask register using EVEX.aaa and EVEX.z

Table 4-11. #UD Conditions of Opmask Related Encoding Field

Notation	Position	Operand Encoding	64-bit #UD	Non-64-bit #UD
EVEX.aaa	P[18 : 16]	instructions do not use opmask for conditional processing ¹	if aaa != 000b	if aaa != 000b
		opmask used as conditional processing mask and updated at completion ²	if aaa = 000b	if aaa = 000b;
		opmask used as conditional processing	None (valid ³)	None (valid ¹)
EVEX.z	P[23]	vector instruction using opmask as source or destination ⁴	if EVEX.z != 0	if EVEX.z != 0
		store instructions or gather/scatter instructions	if EVEX.z != 0	if EVEX.z != 0
		instruction supporting conditional processing mask with EVEX.aaa = 000b	if EVEX.z != 0	if EVEX.z != 0

NOTES:

1. E.g. VBROADCASTMxxx, VPMOVM2x, VPMOVx2M
2. E.g. Gather/Scatter family
3. aaa can take any value. A value of 000 indicates that there is no masking on the instruction; in this case, all elements will be processed as if there was a mask of 'all ones' regardless of the actual value in K0.
4. E.g. VFPCCLASSPD/PS, VCMPPB/D/Q/W family, VPMOVM2x, VPMOVx2M

Table 4-12 lists the #UD conditions of EVEX bit fields that depends on the context of EVEX.b.

Table 4-12. #UD Conditions Dependent on EVEX.b Context

Notation	Position	Operand Encoding	64-bit #UD	Non-64-bit #UD
EVEX.L'Lb	P[22 : 20]	reg-reg, FP instructions with rounding semantic	None (valid ¹)	None (valid ¹)
		other reg-reg, FP instructions that can cause #XF	None (valid ²)	None (valid ²)
		other reg-mem instructions in Table 4-5	None (valid ³)	None (valid ³)
		other instruction classes ⁴ in Table 4-6	if EVEX.b > 0	if EVEX.b > 0

NOTES:

1. L'L specifies rounding control, see Table 4-7, supports {er} syntax.
2. L'L specifies vector length, see Table 4-7, supports {sae} syntax.
3. L'L specifies vector length, see Table 4-7, supports embedded broadcast syntax
4. L'L specifies either vector length or ignored.

4.8 DEVICE NOT AVAILABLE

EVEX-encoded instructions follow the same rules when it comes to generating #NM (Device Not Available) exception. In particular, it is generated when CR0.TS[bit 3] = 1.

4.9 SCALAR INSTRUCTIONS

EVEX-encoded scalar SIMD instructions can access up to 32 registers in 64-bit mode. Scalar instructions support masking (using the least significant bit of the opmask register), but broadcasting is not supported.

4.10 EXCEPTION CLASSIFICATIONS OF EVEX-ENCODED INSTRUCTIONS

The exception behavior of EVEX-encoded instructions can be classified into the classes shown in the rest of this section. The classification of EVEX-encoded instructions follow a similar framework as those of AVX and AVX2 instructions using the VEX prefix. Exception types for EVEX-encoded instructions are named in the style of “E##” or with a suffix “E##XX”. The “##” designation generally follows that of AVX/AVX2 instructions. The majority of EVEX encoded instruction with “Load+op” semantic supports memory fault suppression, which is represented by E##. The instructions with “Load+op” semantic but do not support fault suppression are named “E##NF”. A summary table of exception classes by class names are shown below.

Table 4-13. EVEX-Encoded Instruction Exception Class Summary

Exception Class	Instruction set	Mem arg	(#XM)
Type E1	Vector Moves/Load/Stores	explicitly aligned, w/ fault suppression	none
Type E1NF	Vector Non-temporal Stores	explicitly aligned, no fault suppression	none
Type E2	FP Vector Load+op	Support fault suppression	yes
Type E2NF	FP Vector Load+op	No fault suppression	yes
Type E3	FP Scalar/Partial Vector, Load+Op	Support fault suppression	yes
Type E3NF	FP Scalar/Partial Vector, Load+Op	No fault suppression	yes
Type E4	Integer Vector Load+op	Support fault suppression	no
Type E4NF	Integer Vector Load+op	No fault suppression	no
Type E5	Legacy-like Promotion	Varies, Support fault suppression	no
Type E5NF	Legacy-like Promotion	Varies, No fault suppression	no
Type E6	Post AVX Promotion	Varies, w/ fault suppression	no
Type E6NF	Post AVX Promotion	Varies, no fault suppression	no
Type E7NM	register-to-register op	none	none
Type E9NF	Miscellaneous 128-bit	Vector-length Specific, no fault suppression	none
Type E10	Non-XF Scalar	Vector Length ignored, w/ fault suppression	none
Type E10NF	Non-XF Scalar	Vector Length ignored, no fault suppression	none
Type E11	VCVTPH2PS	Half Vector Length, w/ fault suppression	yes
Type E11NF	VCVTPS2PH	Half Vector Length, no fault suppression	yes
Type E12	Gather and Scatter Family	VSIB addressing, w/ fault suppression	none
Type E12NP	Gather and Scatter Prefetch Family	VSIB addressing, w/o page fault	none

Table 4-14 lists EVEX-encoded instruction mnemonic by exception classes.

Table 4-14. EVEX Instructions in each Exception Class

Exception Class	Instruction
Type E1	VMOVAPD, VMOVAPS, VMOVDQA32, VMOVDQA64
Type E1NF	VMOVNTDQ, VMOVNTDQA, VMOVNTPD, VMOVNTPS
Type E2	VADDPD, VADDPS, VCMPPD, VCMPPS, VCVTDQ2PS, VCVTPD2DQ, VCVTPD2PS, VCVTPS2DQ, VCVTTPD2DQ, VCVTTPS2DQ, VDIVPD, VDIVPS, VFMADDxxxPD, VFMADDxxxPS, VFMSUBADDxxxPD, VFMSUBADDxxxPS, VFMSUBxxxPD, VFMSUBxxxPS, VFNMADDxxxPD, VFNMADDxxxPS, VFNMSUBxxxPD, VFNMSUBxxxPS, VMAXPD, VMAXPS, VMINPD, VMINPS, VMULPD, VMULPS, VSQRTPD, VSQRTPS, VSUBPD, VSUBPS
	VCVTPD2QQ, VCVTPD2UQQ, VCVTPD2UDQ, VCVTPS2UDQS, VCVTQQ2PD, VCVTQQ2PS, VCVTTPD2DQ, VCVTTPD2QQ, VCVTTPD2UDQ, VCVTTPD2UQQ, VCVTTPS2DQ, VCVTTPS2UDQ, VCVTUDQ2PS, VCVTUQQ2PD, VCVTUQQ2PS, VFIXUPIMMPD, VFIXUPIMMPS, VGETEXPPD, VGETEXPPS, VGETMANTPD, VGETMANTPS, VRANGEPD, VRANGEPS, VREDUCEPD, VREDUCEPS, VRNDSCALEPD, VRNDSCALEPS, VSCALEFPD, VSCALEFPS, VRCP28PD, VRCP28PS, VRSQRT28PD, VRSQRT28PS
Type E3	VADDSB, VADDS, VCMPSD, VCMPS, VCVTPS2PD, VCVTSD2SS, VCVTSS2SD, VDIVSD, VDIVSS, VMAXSD, VMAXSS, VMINSD, VMINSS, VMULSD, VMULSS, VSQRTSD, VSQRTSS, VSUBSD, VSUBSS
	VCVTPS2QQ, VCVTPS2UQQ, VCVTTPS2QQ, VCVTTPS2UQQ, VFMADDxxxSD, VFMADDxxxSS, VFMSUBxxxSD, VFMSUBxxxSS, VFNMADDxxxSD, VFNMADDxxxSS, VFNMSUBxxxSD, VFNMSUBxxxSS, VFIXUPIMMSD, VFIXUPIMMSS, VGETEXPSD, VGETEXPSS, VGETMANTSD, VGETMANTSS, VRANGESD, VRANGESS, VREDUCESD, VREDUCESS, VRNDSCALESD, VRNDSCALESS, VSCALEFSD, VSCALEFSS, VRCP28SD, VRCP28SS, VRSQRT28SD, VRSQRT28SS
Type E3NF	VCOMISD, VCOMISS, VCVTSD2SI, VCVTSI2SD, VCVTSI2SS, VCVTSS2SI, VCVTSS2SD, VCVTSS2SI, VUCOMISD, VUCOMISS
	VCVTSD2USI, VCVTSS2USI, VCVTSS2USI, VCVTSS2USI, VCVTUSI2SD, VCVTUSI2SS
Type E4	VANDPD, VANDPS, VANDNPD, VANDNPS, VORPD, VORPS, VPABSD, VPABSQ, VPADD, VPADDQ, VPAND, VPANDQ, VPANDND, VPANDNQ, VPCMPDQ, VPCMPDQ, VPCMPGT, VPCMPGTQ, VPMAXSD, VPMAXSQ, VPMAXUD, VPMAXUQ, VPMINSD, VPMINSQ, VPMINUD, VPMINUQ, VPMULLD, VPMULLQ, VPMULUDQ, VPMULDQ, VPORD, VPORQ, VPSUBD, VPSUBQ, VPXORD, VPXORQ, VXORPD, VXORPS, VPSLLVD, VPSLLVQ,
	VBLENDMPD, VBLENDMPS, VBLENDMD, VBLENDMQ, VFPCLASSPD, VFPCLASSPSP, VPCMPD, VPCMPQ, VPCMPUD, VPCMPUQ, VPLZCNTD, VPLZCNTQ, VPROLD, VPROLQ, (VPSLLD, VPSLLQ, VPSRAD, VPSRAQ, VPSRLD, VPSRLQ) ¹ , VPTERNLOGD, VPTERNLOGQ, VPTESTMD, VPTESTMQ, VPTESTNMD, VPTESTNMQ, VRCP14PD, VRCP14PS, VRSQRT14PD, VRSQRT14PS, VPCONFLICTD, VPCONFLICTQ, VPSRAVW, VPSRAVD, VPSRAVW, VPSRAVQ, VPMADD52LUQ, VPMADD52HUQ
E4.nb ²	VMOVUPD, VMOVUPS, VMOVDQU8, VMOVDQU16, VMOVDQU32, VMOVDQU64, VPCMPB, VPCMPW, VPCMPUB, VPCMPUW, VEXPANDPD, VEXPANDPS, VPCOMPRESSD, VPCOMPRESSQ, VPEXPANDD, VPEXPANDQ, VCOMPRESSPD, VCOMPRESSPS, VPABSB, VPABSW, VPADDB, VPADDW, VPADDSB, VPADDW, VPADDUSB, VPADDUSW, VPAVGB, VPAVGW, VPCMPQB, VPCMPQW, VPCMPGTB, VPCMPGTW, VPMAXSB, VPMAXSW, VPMAXB, VPMAXW, VPMINSB, VPMINSW, VPMINUB, VPMINUW, VPMULHRW, VPMULHW, VPMULLW, VPSUBB, VPSUBW, VPSUBSB, VPSUBSW, VPTESTMB, VPTESTMW, VPTESTNMB, VPTESTNMW, VPSLLW, VPSRAW, VPSRLW, VPSLLVW, VPSRLVW
Type E4NF	VPACKSSDW, VPACKUSDW, VPSHUF, VPUNPCKHDQ, VPUNPCKHQDQ, VPUNPCKLDQ, VPUNPCKLQDQ, VSHUFPD, VSHUFPS, VUNPCKHPD, VUNPCKHPS, VUNPCKLPD, VUNPCKLPS, VPERMD, VPERMPS, VPERMPD, VPERMQ,
	VALIGND, VALIGNQ, VPERMI2D, VPERMI2PS, VPERMI2PD, VPERMI2Q, VPERMT2D, VPERMT2PS, VPERMT2Q, VPERMT2PD, VPERMILPD, VPERMILPS, VSHUFI64X2, VSHUFI64X2, VSHUFF32X4, VSHUFF64X2, VPMULTISHIFTQB
E4NF.nb ²	VDBPSADBW, VPACKSSWB, VPACKUSWB, VPALIGNR, VPMADDWD, VPMADDUBSW, VMOVSHDUP, VMOVSLDUP, VPSADBW, VPSHUF, VPSHUFHW, VPSHUFHW, VPSLLDQ, VPSRLDQ, VPSLLW, VPSRAW, VPSRLW, (VPSLLD, VPSLLQ, VPSRAD, VPSRAQ, VPSRLD, VPSRLQ) ³ , VPUNPCKHBW, VPUNPCKHWD, VPUNPCKLBW, VPUNPCKLWD, VPERMW, VPERMI2W, VPERMT2W, VPERMB, VPERMI2B, VPERMT2B

Table 4-14. EVEX Instructions in each Exception Class(Continued)

Exception Class	Instruction
Type E5	VCVTDQ2PD, PMOVSWBW, PMOVSWBW, PMOVSWBD, PMOVSWBQ, PMOVSWWD, PMOVSWWQ, PMOVSDQ, PMOVZXBW, PMOVZXBW, PMOVZXBQ, PMOVZXWD, PMOVZXWQ, PMOVZXDQ
	VCVTUDQ2PD
Type E5NF	VMOVDDUP
Type E6	VBROADCASTSS, VBROADCASTSD, VBROADCASTF32X4, VBROADCASTI32X4, VPBROADCASTB, VPBROADCASTD, VPBROADCASTW, VPBROADCASTQ, VBROADCASTF32X2, VBROADCASTF32X4, VBROADCASTF64X2, VBROADCASTF32X8, VBROADCASTF64X4, VBROADCASTI32X2, VBROADCASTI32X4, VBROADCASTI64X2, VBROADCASTI32X8, VBROADCASTI64X4, VFPCLASSSD, VFPCLASSSS, VPMOVQB, VPMOVQB, VPMOVUSQB, VPMOVQW, VPMOVQW, VPMOVUSQW, VPMOVQD, VPMOVQD, VPMOVUSQD, VPMOVDB, VPMOVDB, VPMOVUSDB, VPMOVDW, VPMOVSDW, VPMOVUSDW
Type E6NF	VEXTRACTF32X4, VEXTRACTF64X2, VEXTRACTF32X8, VINSERTF32X4, VINSERTF64X2, VINSERTF64X4, VINSERTF32X8, VINSERTI32X4, VINSERTI64X2, VINSERTI64X4, VINSERTI32X8, VEXTRACTI32X4, VEXTRACTI64X2, VEXTRACTI32X8, VEXTRACTI64X4, VPBROADCASTMB2Q, VPBROADCASTMW2D, VPMOVWB, VPMOVSWB, VPMOVUSWB
Type E7NM.128 ⁴	VMOVLHPS, VMOVHPS
Type E7NM.	(VPBROADCASTD, VPBROADCASTQ, VPBROADCASTB, VPBROADCASTW) ⁵ , VPMOVM2B, VPMOVM2D, VPMOVM2Q, VPMOVM2W, VPMOVMB2M, VPMOVM2M, VPMOVQ2M, VPMOVW2M
Type E9NF	VEXTRACTPS, VINSERTPS, VMOVHPD, VMOVHPS, VMOVLPD, VMOVLPS, VMOVD, VMOVQ, VPEXTRB, VPEXTRD, VPEXTRW, VPEXTRQ, VPINSRB, VPINSRD, VPINSRW, VPINSRQ
Type E10	VMOVSD, VMOVSS, VRCPP14SD, VRCPP14SS, VRSQRT14SD, VRSQRT14SS,
Type E10NF	(VCVTISI2SD, VCVTUSI2SD) ⁶
Type E11	VCVTPH2PS, VCVTPS2PH
Type E12	VGATHERDPS, VGATHERDPD, VGATHERQPS, VGATHERQPD, VPGATHERDD, VPGATHERDQ, VPGATHERQD, VPGATHERQQ, VPSCATTERDD, VPSCATTERDQ, VPSCATTERQD, VPSCATTERQQ, VSCATTERDPD, VSCATTERDPS, VSCATTERQPD, VSCATTERQPS
Type E12NP	VGATHERPFODPD, VGATHERPFODPS, VGATHERPFOQPD, VGATHERPFOQPS, VGATHERPF1DPD, VGATHERPF1DPS, VGATHERPF1QPD, VGATHERPF1QPS, VSCATTERPFODPD, VSCATTERPFODPS, VSCATTERPFOQPD, VSCATTERPFOQPS, VSCATTERPF1DPD, VSCATTERPF1DPS, VSCATTERPF1QPD, VSCATTERPF1QPS

NOTES:

1. Operand encoding FVI tupletype with immediate.
2. Embedded broadcast is not supported with the “.nb” suffix.
3. Operand encoding M128 tupletype.
4. #UD raised if EVEX.L'L !=00b (VL=128).
5. The source operand is a general purpose register.
6. W0 encoding only.

4.10.1 Exceptions Type E1 and E1NF of EVEX-Encoded Instructions

EVEX-encoded instructions with memory alignment restrictions, and supporting memory fault suppression follow exception class E1.

Table 4-15. Type E1 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X	X	EVEX.512: Memory operand is not 64-byte aligned. EVEX.256: Memory operand is not 32-byte aligned. EVEX.128: Memory operand is not 16-byte aligned.
			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	If fault suppression not set, and a page fault.

EVEX-encoded instructions with memory alignment restrictions, but do not support memory fault suppression follow exception class E1NF.

Table 4-16. Type E1NF Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X	X	EVEX.512: Memory operand is not 64-byte aligned. EVEX.256: Memory operand is not 32-byte aligned. EVEX.128: Memory operand is not 16-byte aligned.
			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	For a page fault.

4.10.2 Exceptions Type E2 of EVEX-Encoded Instructions

EVEX-encoded vector instructions with arithmetic semantic follow exception class E2.

Table 4-17. Type E2 Class Exception Conditions

Exception	Real	Virtual 8086	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
	X	X	X	X	If an unmasked SIMD floating-point exception and CR4.OSXMMEXCPT[bit 10] = 0.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> ▪ State requirement, Table 4-8 not met. ▪ Opcode independent #UD condition in Table 4-9. ▪ Operand encoding #UD conditions in Table 4-10. ▪ Opmask encoding #UD condition of Table 4-11. ▪ If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	If fault suppression not set, and a page fault.
SIMD Floating-point Exception, #XM	X	X	X	X	If an unmasked SIMD floating-point exception, {sae} or {er} not set, and CR4.OSXMMEXCPT[bit 10] = 1.

4.10.3 Exceptions Type E3 and E3NF of EVEX-Encoded Instructions

EVEX-encoded scalar instructions with arithmetic semantic that support memory fault suppression follow exception class E3.

Table 4-18. Type E3 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
	X	X	X	X	If an unmasked SIMD floating-point exception and CR4.OSXMMEXCPT[bit 10] = 0.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0.
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	If fault suppression not set, and a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.
SIMD Floating-point Exception, #XM	X	X	X	X	If an unmasked SIMD floating-point exception, {sae} or {er} not set, and CR4.OSXMMEXCPT[bit 10] = 1.

EVEX-encoded scalar instructions with arithmetic semantic that do not support memory fault suppression follow exception class E3NF.

Table 4-19. Type E3NF Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			EVEX prefix.
	X	X	X	X	If an unmasked SIMD floating-point exception and CR4.OSXMMEXCPT[bit 10] = 0.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0.
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	For a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.
SIMD Floating-point Exception, #XM	X	X	X	X	If an unmasked SIMD floating-point exception, {sae} or {er} not set, and CR4.OSXMMEXCPT[bit 10] = 1.

4.10.4 Exceptions Type E4 and E4NF of EVEX-Encoded Instructions

EVEX-encoded vector instructions that cause no SIMD FP exception and support memory fault suppression follow exception class E4.

Table 4-20. Type E4 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0 and in E4.nb subclass (see E4.nb entries in Table 4-14). If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	If fault suppression not set, and a page fault.

EVEX-encoded vector instructions that do not cause SIMD FP exception nor support memory fault suppression follow exception class E4NF.

Table 4-21. Type E4NF Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0 and in E4NF.nb subclass (see E4NF.nb entries in Table 4-14). If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	For a page fault.

4.10.5 Exceptions Type E5 and E5NF

EVEX-encoded scalar/partial-vector instructions that cause no SIMD FP exception and support memory fault suppression follow exception class E5.

Table 4-22. Type E5 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (FOH).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	If fault suppression not set, and a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.

EVEX-encoded scalar/partial vector instructions that do not cause SIMD FP exception nor support memory fault suppression follow exception class E5NF.

Table 4-23. Type E5NF Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	For a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.

4.10.6 Exceptions Type E6 and E6NF

Table 4-24. Type E6 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512).
			X	X	If preceded by a LOCK prefix (FOH).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
			X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM			X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
Page Fault #PF(fault-code)			X	X	If fault suppression not set, and a page fault.
Alignment Check #AC(0)			X	X	For 4 or 8 byte memory references if alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.

EVEX-encoded instructions that do not cause SIMD FP exception nor support memory fault suppression follow exception class E6NF.

Table 4-25. Type E6NF Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512).
			X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
			X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM			X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If the memory address is in a non-canonical form.
Page Fault #PF(fault-code)			X	X	For a page fault.
Alignment Check #AC(0)			X	X	For 4 or 8 byte memory references if alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.

4.10.7 Exceptions Type E7NM

EVEX-encoded instructions that cause no SIMD FP exception and do not reference memory follow exception class E7NM.

Table 4-26. Type E7NM Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> ▪ State requirement, Table 4-8 not met. ▪ Opcode independent #UD condition in Table 4-9. ▪ Operand encoding #UD conditions in Table 4-10. ▪ Opmask encoding #UD condition of Table 4-11. ▪ If EVEX.b != 0. ▪ Instruction specific EVEX.L'L restriction not met.
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM			X	X	If CR0.TS[bit 3]=1.

4.10.8 Exceptions Type E9 and E9NF

EVEX-encoded vector or partial-vector instructions that do not cause no SIMD FP exception and support memory fault suppression follow exception class E9.

Table 4-27. Type E9 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 00b (VL=128).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	If fault suppression not set, and a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.

EVEX-encoded vector or partial-vector instructions that must be encoded with VEX.L'L = 0, do not cause SIMD FP exception nor support memory fault suppression follow exception class E9NF.

Table 4-28. Type E9NF Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 00b (VL=128).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	For a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference is made while the current privilege level is 3.

4.10.9 Exceptions Type E10

EVEX-encoded scalar instructions that ignore EVEX.L'L vector length encoding and do not cause no SIMD FP exception, support memory fault suppression follow exception class E10.

Table 4-29. Type E10 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0.
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	If fault suppression not set, and a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.

EVEX-encoded scalar instructions that must be encoded with VEX.L'L = 0, do not cause SIMD FP exception nor support memory fault suppression follow exception class E10NF.

Table 4-30. Type E10NF Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0.
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	If fault suppression not set, and a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.

4.10.10 Exception Type E11 (EVEX-only, mem arg no AC, floating-point exceptions)

EVEX-encoded instructions that can cause SIMD FP exception, memory operand support fault suppression but do not cause #AC follow exception class E11.

Table 4-31. Type E11 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a EVEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF (fault-code)		X	X	X	If fault suppression not set, and a page fault.
SIMD Floating-Point Exception, #XM	X	X	X	X	If an unmasked SIMD floating-point exception, {sae} not set, and CR4.OSXMMEX-CPT[bit 10] = 1.

4.10.11 Exception Type E12 and E12NP (VSIB mem arg, no AC, no floating-point exceptions)

Table 4-32. Type E12 Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512). If vvvv != 1111b.
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a VEX prefix.
	X	X	X	NA	If address size attribute is 16 bit.
	X	X	X	X	If ModR/M.mod = '11b'.
	X	X	X	X	If ModR/M.rm != '100b'.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
	X	X	X	X	If k0 is used (gather or scatter operation).
	X	X	X	X	If index = destination register (gather operation).
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF (fault-code)		X	X	X	For a page fault.

EVEX-encoded prefetch instructions that do not cause #PF follow exception class E12NP.

Table 4-33. Type E12NP Class Exception Conditions

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X			If EVEX prefix present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10. Opmask encoding #UD condition of Table 4-11. If EVEX.b != 0. If EVEX.L'L != 10b (VL=512).
	X	X	X	X	If preceded by a LOCK prefix (F0H).
			X	X	If any REX, F2, F3, or 66 prefixes precede a VEX prefix.
	X	X	X	NA	If address size attribute is 16 bit.
	X	X	X	X	If ModR/M.mod = '11b'.
	X	X	X	X	If ModR/M.rm != '100b'.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
	X	X	X	X	If k0 is used (gather or scatter operation).
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.

4.11 EXCEPTION CLASSIFICATIONS OF OPMASK INSTRUCTIONS

The exception behavior of VEX-encoded opmask instructions are listed below.

Exception conditions of Opmask instructions that do not address memory are listed as Type K20.

Table 4-34. TYPE K20 Exception Definition (VEX-Encoded OpMask Instructions w/o Memory Arg)

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X	X	X	If relevant CPUID feature flag is '0'.
	X	X			If a VEX prefix is present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10.
			X	X	If any REX, F2, F3, or 66 prefixes precede a VEX prefix.
			X	X	If ModRM:[7:6] != 11b.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.

Exception conditions of Opmask instructions that address memory are listed as Type K21.

Table 4-35. TYPE K21 Exception Definition (VEX-Encoded OpMask Instructions Addressing Memory)

Exception	Real	Virtual 80x86	Protected and Compatibility	64-bit	Cause of Exception
Invalid Opcode, #UD	X	X	X	X	If relevant CPUID feature flag is '0'.
	X	X			If a VEX prefix is present.
			X	X	If CR4.OSXSAVE[bit 18]=0. If any one of following conditions applies: <ul style="list-style-type: none"> State requirement, Table 4-8 not met. Opcode independent #UD condition in Table 4-9. Operand encoding #UD conditions in Table 4-10.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
			X	X	If any REX, F2, F3, or 66 prefixes precede a VEX prefix.
Stack, SS(0)	X	X	X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If the DS, ES, FS, or GS register is used to access memory and it contains a null segment selector.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault #PF(fault-code)		X	X	X	For a page fault.
Alignment Check #AC(0)		X	X	X	If alignment checking is enabled and an unaligned memory reference of 8 bytes or less is made while the current privilege level is 3.

CHAPTER 5

INSTRUCTION SET REFERENCE, A-Z

Instructions described in this document follow the general documentation convention established in *Intel 64 and IA-32 Architectures Software Developer's Manual Volume 2A*. Additional notations and conventions adopted in this document are listed in *Section 5.1*. *Section 5.1.5.1* covers supplemental information that applies to a specific subset of instructions.

5.1 INTERPRETING INSTRUCTION REFERENCE PAGES

This section describes the format of information contained in the instruction reference pages in this chapter. It explains notational conventions and abbreviations used in these sections that are outside of those conventions described in *Section 3.1* of the *Intel 64 and IA-32 Architectures Software Developer's Manual Volume 2A*.

5.1.1 Instruction Format

The following is an example of the format used for each instruction description in this chapter. The table below provides an example summary table.

ADDPS—Add Packed Single-Precision Floating-Point Values (THIS IS AN EXAMPLE)

Opcode/ Instruction	64/32 bit Mode Support	CPUID Feature Flag	Description
0F 5B /r ADDPS xmm1, xmm2/m128	V/V	SSE	Add packed single-precision floating-point values from xmm2/mem to xmm1 and store result in xmm1.
VEX.NDS.128.0F 5B /r VADDPS xmm1, xmm2, xmm3/m128	V/V	AVX	Add packed single-precision floating-point values from xmm3/mem to xmm2 and store result in xmm1.
VEX.NDS.256.0F 5B /r VADDPS ymm1, ymm2, ymm3/m256	V/V	AVX	Add packed single-precision floating-point values from ymm3/mem to ymm2 and store result in ymm1.
VEX.L1.0F.W0 41 /r KANDW k1, k2, k3	V/V	AVX512F	Bitwise AND word masks k2 and k3 and place result in k1.
EVEX.NDS.128.0F.W0 5B /r VADDPS xmm1 {k1}{z}, xmm2, xmm3/m128/m32bcst	V/V	AVX512VL AVX512F	Add packed single-precision floating-point values from xmm3/m128/m32bcst to xmm2 and store result in xmm1 with writemask k1.
EVEX.NDS.256.0F.W0 5B /r VADDPS ymm1 {k1}{z}, ymm2, ymm3/m256/m32bcst	V/V	AVX512VL AVX512F	Add packed single-precision floating-point values from ymm3/m256/m32bcst to ymm2 and store result in ymm1 with writemask k1.
EVEX.NDS.512.0F.W0 5B /r VADDPS zmm1 {k1}{z}, zmm2, zmm3/m512/m32bcst {er}	V/V	AVX512F	Add packed single-precision floating-point values from zmm3/m512/m32bcst with zmm2 and store result in zmm1 with writemask k1.

5.1.2 Opcode Column in the Instruction Summary Table

For notation and conventions applicable to instructions that do not use VEX or EVEX prefixes, consult *Section 3.1* of the *Intel 64 and IA-32 Architectures Software Developer's Manual Volume 2A*.

In the Instruction Summary Table, the Opcode column presents each instruction encoded using the VEX prefix in following form (including the modR/M byte if applicable, the immediate byte if applicable):

**VEX.[NDS/NDD/DS].[128,256,LO,L1,LIG].[66,F2,F3].0F/OF3A/OF38.[W0,W1,WIG] opcode [/r]
[ib,/is4]**

- **VEX:** indicates the presence of the VEX prefix is required. The VEX prefix can be encoded using the three-byte form (the first byte is C4H), or using the two-byte form (the first byte is C5H). The two-byte form of VEX only applies to those instructions that do not require the following fields to be encoded: VEX.mmmmm, VEX.W, VEX.X, VEX.B.

The encoding of various sub-fields of the VEX prefix is described using the following notations:

- **NDS, NDD, DDS:** implies that VEX.vvvv field is valid for the encoding of an operand. It may specify either the source register (NDS) or the destination register (NDD). The VEX.vvvv field can be encoded using either the 2-byte or 3-byte form of the VEX prefix. DDS expresses a syntax where vvvv encodes the second source register in a three-operand instruction syntax where the content of first source register will be overwritten by the result. If NDS, NDD and DDS are absent (i.e. VEX.vvvv does not encode an operand), VEX.vvvv must be 1111b.
- **128,256,L0,L1:** VEX.L fields can be 0 (denoted by VEX.128 or VEX.L0 for mask instructions) or 1 (denoted by VEX.256 or VEX.L1 for mask instructions). The VEX.L field can be encoded using either the 2-byte or 3-byte form of the VEX prefix. The presence of the notation VEX.256 or VEX.128 in the opcode column should be interpreted as follows:
 - If VEX.256 is present in the opcode column: The semantics of the instruction must be encoded with VEX.L = 1. An attempt to encode this instruction with VEX.L = 0 can result in one of two situations: (a) if VEX.128 version is defined, the processor will behave according to the defined VEX.128 behavior; (b) an #UD occurs if there is no VEX.128 version defined.
 - If VEX.128 is present in the opcode column but there is no VEX.256 version defined for the same opcode byte: Three situations apply: (a) For VEX-encoded, 128-bit SIMD integer instructions, software must encode the instruction with VEX.L = 0. The processor will treat the opcode byte encoded with VEX.L = 1 by causing an #UD exception; (b) For VEX-encoded, 128-bit packed floating-point instructions, software must encode the instruction with VEX.L = 0. The processor will treat the opcode byte encoded with VEX.L = 1 by causing an #UD exception (e.g. VMOVLPS); (c) For VEX-encoded, scalar, SIMD floating-point instructions, software should encode the instruction with VEX.L = 0 to ensure software compatibility with future processor generations. Scalar SIMD floating-point instruction can be distinguished from the mnemonic of the instruction. Generally, the last two letters of the instruction mnemonic would be either "SS", "SD", or "SI" for SIMD floating-point conversion instructions, except VBROADCASTSx are unique cases.
 - VEX.L0 and VEX.L1 notations are used in the case of masking instructions such as KANDW since the VEX.L bit is not used to distinguish between the 128-bit and 256-bit forms for these instructions. Instead, this bit is used to distinguish between the two operand form (VEX.L0) and the three operand form (VEX.L1) of the same mask instruction.
 - If VEX.L0 is present in the opcode column: The semantics of the instruction must be encoded with VEX.L = 0. An attempt to encode this instruction with VEX.L = 1 can result in one of two situations: (a) if VEX.L1 version is defined, the processor will behave according to the defined VEX.L1 behavior; (b) an #UD occurs if there is no VEX.L1 version defined.
 - If VEX.L1 is present in the opcode column: The semantics of the instruction must be encoded with VEX.L = 1. An attempt to encode this instruction with VEX.L = 0 can result in one of two situations: (a) if VEX.L0 version is defined, the processor will behave according to the defined VEX.L1 behavior; (b) an #UD occurs if there is no VEX.L0 version defined.
 - **LIG:** VEX.L bit ignored
- **66,F2,F3:** The presence or absence of these value maps to the VEX.pp field encodings. If absent, this corresponds to VEX.pp=00B. If present, the corresponding VEX.pp value affects the "opcode" byte in the same way as if a SIMD prefix (66H, F2H or F3H) does to the ensuing opcode byte. Thus a non-zero encoding of VEX.pp may be considered as an implied 66H/F2H/F3H prefix. The VEX.pp field may be encoded using either the 2-byte or 3-byte form of the VEX prefix.
- **0F,0F3A,0F38:** The presence maps to a valid encoding of the VEX.mmmmm field. Only three encoded values of VEX.mmmmm are defined as valid, corresponding to the escape byte sequence of 0FH, 0F3AH and 0F38H. The effect of a valid VEX.mmmmm encoding on the ensuing opcode byte is the same as if the corresponding escape byte sequence on the ensuing opcode byte for non-VEX encoded instructions. Thus a valid encoding of VEX.mmmmm may be considered as an implied escape byte sequence of either 0FH, 0F3AH or 0F38H. The VEX.mmmmm field must be encoded using the 3-byte form of VEX prefix.

- **0F,0F3A,0F38 and 2-byte/3-byte VEX.** The presence of 0F3A and 0F38 in the opcode column implies that opcode can only be encoded by the three-byte form of VEX. The presence of 0F in the opcode column does not preclude the opcode to be encoded by the two-byte of VEX if the semantics of the opcode does not require any subfield of VEX not present in the two-byte form of the VEX prefix.
 - **W0:** VEX.W=0.
 - **W1:** VEX.W=1.
 - **WIG:** VEX.W bit ignored
 - The presence of W0/W1 in the opcode column applies to two situations: (a) it is treated as an extended opcode bit, (b) the instruction semantics support an operand size promotion to 64-bit of a general-purpose register operand or a 32-bit memory operand. The presence of W1 in the opcode column implies the opcode must be encoded using the 3-byte form of the VEX prefix. The presence of W0 in the opcode column does not preclude the opcode to be encoded using the C5H form of the VEX prefix, if the semantics of the opcode does not require other VEX subfields not present in the two-byte form of the VEX prefix. If neither W0 or W1 is present, the instruction may be encoded using either the two-byte form (if the opcode semantic does not require VEX subfields not present in the two-byte form of VEX) or the three-byte form of VEX. Encoding an instruction using the two-byte form of VEX is equivalent to W0.
 - **opcode:** Instruction opcode.
 - **ib:** An 8-bit immediate byte is present and used as one of the instructions operands.
 - **/is4:** An 8-bit immediate byte is present containing a source register specifier in imm[7:4] and instruction-specific payload in imm[3:0].
 - **imz2:** Part of the is4 immediate byte provides control functions that apply to two-source permute instructions
- In general, the encoding of VEX.R, VEX.X, VEX.B field are not shown explicitly in the opcode column.

**EVEX.[NDS/NDD/DDS].[128,256,512,LIG].[66,F2,F3].0F/0F3A/0F38.[W0,W1,WIG] opcode [/r]
[ib,/is4]**

- **EVEX:** The EVEX prefix is encoded using the four-byte form (the first byte is 62H). Refer to Section 4.2 for more detail on the EVEX prefix.
- The encoding of various sub-fields of the EVEX prefix is described using the following notations.
- **NDS, NDD, DDS:** implies that EVEX.vvvv (and EVEX.v') field is valid for the encoding of an operand. It may specify either the source register (NDS) or the destination register (NDD). DDS expresses a syntax where vvvv encodes the second source register in a three-operand instruction syntax where the content of first source register will be overwritten by the result. If both NDS and NDD absent (i.e. EVEX.vvvv does not encode an operand), EVEX.vvvv must be 1111b (and EVEX.v' must be 1b).
 - **128, 256, 512, LIG:** This corresponds to the vector length; three values are allowed by EVEX: 512-bit, 256-bit and 128-bit. Alternatively, vector length is ignored (LIG) for certain instructions; this typically applies to scalar instructions operating on one data element of a vector register.
 - **66,F2,F3:** The presence of these value maps to the EVEX.pp field encodings. The corresponding VEX.pp value affects the "opcode" byte in the same way as if a SIMD prefix (66H, F2H or F3H) does to the ensuing opcode byte. Thus a non-zero encoding of VEX.pp may be considered as an implied 66H/F2H/F3H prefix.
 - **0F,0F3A,0F38:** The presence maps to a valid encoding of the EVEX.mmm field. Only three encoded values of EVEX.mmm are defined as valid, corresponding to the escape byte sequence of 0FH, 0F3AH and 0F38H. The effect of a valid EVEX.mmm encoding on the ensuing opcode byte is the same as if the corresponding escape byte sequence on the ensuing opcode byte for non-EVEX encoded instructions. Thus a valid encoding of EVEX.mmm may be considered as an implied escape byte sequence of either 0FH, 0F3AH or 0F38H.
 - **W0:** EVEX.W=0.
 - **W1:** EVEX.W=1.
 - **WIG:** EVEX.W bit ignored
 - **opcode:** Instruction opcode.

- **/is4**: An 8-bit immediate byte is present containing a source register specifier in imm[7:4] and instruction-specific payload in imm[3:0].
- **imz2**: Part of the is4 immediate byte provides control functions that apply to two-source permute instructions
- In general, the encoding of EVEX.R and R', EVEX.X and X', and EVEX.B and B' fields are not shown explicitly in the opcode column.

5.1.3 Instruction Column in the Instruction Summary Table

- **xmm** — an XMM register. The XMM registers are: XMM0 through XMM7; XMM8 through XMM15 are available in 64-bit mode. XMM16 through XMM31 are available in 64-bit mode via EVEX prefix.
- **ymm** — a YMM register. The 256-bit YMM registers are: YMM0 through YMM7; YMM8 through YMM15 are available in 64-bit mode. YMM16 through YMM31 are available in 64-bit mode via EVEX prefix.
- **m256** — A 32-byte operand in memory.
- **ymm/m256** - a YMM register or 256-bit memory operand.
- **<YMM0>**: indicates use of the YMM0 register as an implicit argument.
- **zmm** — a ZMM register. The 512-bit ZMM registers require EVEX prefix and are: ZMM0 through ZMM7; ZMM8 through ZMM31 are available in 64-bit mode.
- **m512** — A 64-byte operand in memory.
- **zmm/m512** — a ZMM register or 512-bit memory operand.
- **{k1}{z}** — a mask register used as instruction writemask. The 64-bit k registers are: k1 through k7. Writemask specification is available exclusively via EVEX prefix. The masking can either be done as a merging-masking, where the old values are preserved for masked out elements or as a zeroing masking. The type of masking is determined by using the EVEX.z bit.
- **{k1}** — without {z}: a mask register used as instruction writemask for instructions that do not allow zeroing-masking but support merging-masking. This corresponds to instructions that require the value of the aaa field to be different than 0 (e.g., gather) and store-type instructions which allow only merging-masking.
- **k1** — a mask register used as a regular operand (either destination or source). The 64-bit k registers are: k0 through k7.
- **mV** — a vector memory operand; the operand size is dependent on the instruction.
- **vm32{x,y,z}** — A vector array of memory operands specified using VSIB memory addressing. The array of memory addresses are specified using a common base register, a constant scale factor, and a vector index register with individual elements of 32-bit index value in an XMM register (vm32x), a YMM register (vm32y) or a ZMM register (vm32z).
- **vm64{x,y,z}** — A vector array of memory operands specified using VSIB memory addressing. The array of memory addresses are specified using a common base register, a constant scale factor, and a vector index register with individual elements of 64-bit index value in an XMM register (vm64x), a YMM register (vm64y) or a ZMM register (vm64z).
- **zmm/m512/m32bcst** — an operand that can be a ZMM register, a 512-bit memory location or a 512-bit vector loaded from a 32-bit memory location.
- **zmm/m512/m64bcst** — an operand that can be a ZMM register, a 512-bit memory location or a 512-bit vector loaded from a 64-bit memory location.
- **<ZMM0>** — indicates use of the ZMM0 register as an implicit argument.
- **{er}** indicates support for embedded rounding control, which is only applicable to the register-register form of the instruction. This also implies support for SAE (Suppress All Exceptions).
- **{sae}** indicates support for SAE (Suppress All Exceptions). This is used for instructions that support SAE, but do not support embedded rounding control.
- **SRC1** — Denotes the first source operand in the instruction syntax of an instruction encoded with the EVEX prefix and having two or more source operands.
- **SRC2** — Denotes the second source operand in the instruction syntax of an instruction encoded with the EVEX prefix and having two or more source operands.

- **SRC3** — Denotes the third source operand in the instruction syntax of an instruction encoded with the EVEX prefix and having three source operands.
- **SRC** — The source in a single-source instruction.
- **DST** — the destination in an instruction. This field is encoded by `reg_field`.

5.1.4 64/32 bit Mode Support column in the Instruction Summary Table

The “64/32 bit Mode Support” column in the Instruction Summary table indicates whether an opcode sequence is supported in 64-bit or the Compatibility/other IA32 modes.

The 64-bit mode support is to the left of the ‘slash’ and has the following notation.

- **V** — Supported.
- **I** — Not supported.
- **N.E.** — Indicates an instruction syntax is not encodable in 64-bit mode (it may represent part of a sequence of valid instructions in other modes).
- **N.P.** — Indicates the REX prefix does not affect the legacy instruction in 64-bit mode.
- **N.I.** — Indicates the opcode is treated as a new instruction in 64-bit mode.
- **N.S.** — Indicates an instruction syntax that requires an address override prefix in 64-bit mode and is not supported. Using an address override prefix in 64-bit mode may result in model-specific execution behavior.

The compatibility/Legacy mode support is to the right of the ‘slash’ and has the following notation.

- **V** — Supported.
- **I** — Not supported.
- **N.E.** — Indicates an Intel 64 instruction mnemonics/syntax that is not encodable; the opcode sequence is not applicable as an individual instruction in compatibility mode or IA-32 mode. The opcode may represent a valid sequence of legacy IA-32 instructions.

5.1.5 CPUID Support column in the Instruction Summary Table

The fourth column holds abbreviated CPUID feature flags (e.g. appropriate bits in CPUID.1:ECX, CPUID.1:EDX for SSE/SSE2/SSE3/SSSE3/SSE4.1/SSE4.2/AVX/F16C support; bits in CPUID.(EAX=07H,ECX=0):BCX for AVX2/AVX512F etc) that indicate processor support for the instruction. If the corresponding flag is ‘0’, the instruction will #UD.

For entries that reference to CPUID feature flags listed in Table 2-1, software should follow the detection procedure described in Section 2.1 and Section 2.2.

For entries that reference to CPUID feature flags listed in Table 2-1 and AVX512VL, software should follow the detection procedure described in Section 2.3.

5.1.5.1 Operand Encoding Column in the Instruction Summary Table

The “operand encoding” column is abbreviated as Op/En in the Instruction Summary table heading. Instruction operand encoding information is provided for each assembly instruction syntax using a letter to cross reference to a row entry in the operand encoding definition table that follows the instruction summary table. The operand encoding table in each instruction reference page lists each instruction operand (according to each instruction syntax and operand ordering shown in the instruction column) relative to the ModRM byte, VEX.vvvv field or additional operand encoding placement.

EVEX encoded instructions employ compressed $\text{disp8} \times N$ encoding of the displacement bytes, where N is defined in Table 4-5 and Table 4-6, according to tuple types. The Op/En column of an EVEX encoded instruction uses an abbreviation that corresponds to the tuple type abbreviation (and may include an additional abbreviation related to ModR/M and vvvv encoding). Most EVEX encoded instructions with VEX encoded equivalent have the ModR/M and vvvv encoding order. In such cases, the Tuple abbreviation is shown and the ModR/M, vvvv encoding abbreviation may be omitted.

NOTES

The letters in the Op/En column of an instruction apply **ONLY** to the encoding definition table immediately following the instruction summary table.

In the encoding definition table, the letter ‘r’ within a pair of parentheses denotes the content of the operand will be read by the processor. The letter ‘w’ within a pair of parenthesis denotes the content of the operand will be updated by the processor.

5.2 SUMMARY OF TERMS

- **“Legacy SSE”** — Refers to SSE, SSE2, SSE3, SSSE3, SSE4, and any future instruction sets referencing XMM registers and encoded without a VEX or EVEX prefix.
- **XGETBV, XSETBV, XSAVE, XRSTOR** are defined in *Intel 64 and IA-32 Architectures Software Developer’s Manual, Volumes 3A and Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2C*.
- **VEX** — Refers to a two-byte or three-byte prefix. AVX and FMA instructions are encoded using a VEX prefix.
- **EVEX** — Refers to a four-byte prefix. AVX512F instructions are encoded using an EVEX prefix.
- **VEX.vvvv** — The VEX bit field specifying a source or destination register (in 1’s complement form).
- **rm_field** — Shorthand for the ModR/M *r/m* field and any REX.B
- **reg_field** — Shorthand for the ModR/M *reg* field and any REX.R

5.3 TERNARY BIT VECTOR LOGIC TABLE

VPTERNLOGD/VPTERNLOGQ instructions operate on dword/qword elements and take three bit vectors of the respective input data elements to form a set of 32/64 indices, where each 3-bit value provides an index into an 8-bit lookup table represented by the imm8 byte of the instruction. The 256 possible values of the imm8 byte is constructed as a 16x16 boolean logic table. The 16 rows of the table uses the lower 4 bits of imm8 as row index. The 16 columns are referenced by imm8[7:4]. The 16 columns of the table are present in two halves, with 8 columns shown in Table 5-1 for the column index value between 0:7, followed by Table 5-2 showing the 8 columns corresponding to column index 8:15. This section presents the two-halves of the 256-entry table using a shorthand notation representing simple or compound boolean logic expressions with three input bit source data.

The three input bit source data will be denoted with the capital letters: A, B, C; where A represents a bit from the first source operand (also the destination operand), B and C represent a bit from the 2nd and 3rd source operands.

Each map entry takes the form of a logic expression consisting of one or more component expressions. Each component expression consists of either a unary or binary boolean operator and associated operands. Each binary boolean operator is expressed in lowercase letters, and operands concatenated after the logic operator. The unary operator ‘not’ is expressed using ‘!’. Additionally, the conditional expression “A?B:C” expresses a result returning B if A is set, returning C otherwise.

A binary boolean operator is followed by two operands, e.g. andAB. For a compound binary expression that contain commutative components and comprising the same logic operator, the 2nd logic operator is omitted and three operands can be concatenated in sequence, e.g. andABC. When the 2nd operand of the first binary boolean expression comes from the result of another boolean expression, the 2nd boolean expression is concatenated after the uppercase operand of the first logic expression, e.g. norBnandAC. When the result is independent of an operand, that operand is omitted in the logic expression, e.g. zeros or norCB.

The 3-input expression “majorABC” returns 0 if two or more input bits are 0, returns 1 if two or more input bits are 1. The 3-input expression “minorABC” returns 1 if two or more input bits are 0, returns 0 if two or more input bits are 1.

The building-block bit logic functions used in Table 5-1 and Table 5-2 include:

- Constants: TRUE (1), FALSE (0)
- Unary function: Not (!)
- Binary functions: and, nand, or, nor, xor, xnor
- Conditional function: Select (?:)
- Tertiary functions: major, minor

Table 5-1. Low 8 columns of the 16x16 Map of VPTERNLOG Boolean Logic Operations

Imm	[7:4]							
[3:0]	0H	1H	2H	3H	4H	5H	6H	7H
00H	FALSE	andAnorBC	norBnandAC	andA!B	norCnandBA	andA!C	andAxorBC	andAnandBC
01H	norABC	norCB	norBxorAC	A?!B:norBC	norCxorBA	A?!C:norBC	A?xorBC:norBC	A?nandBC:norBC
02H	andCnorBA	norBxnorAC	andC!B	norBnorAC	C?norBA:andBA	C?norBA:A	C?!B:andBA	C?!B:A
03H	norBA	norBandAC	C?!B:norBA	!B	C?norBA:xnorBA	A?!C:!B	A?xorBC:!B	A?nandBC:!B
04H	andBnorAC	norCxnorBA	B?norAC:andAC	B?norAC:A	andB!C	norCnorBA	B?!C:andAC	B?!C:A
05H	norCA	norCandBA	B?norAC:xnorAC	A?!B:!C	B?!C:norAC	!C	A?xorBC:!C	A?nandBC:!C
06H	norAxnorBC	A?norBC:xorBC	B?norAC:C	xorBorAC	C?norBA:B	xorCorBA	xorCB	B?!C:orAC
07H	norAandBC	minorABC	C?!B:!A	nandBorAC	B?!C:!A	nandCorBA	A?xorBC:nandBC	nandCB
08H	norAnandBC	A?norBC:andBC	andCxorBA	A?!B:andBC	andBxorAC	A?!C:andBC	A?xorBC:andBC	xorAandBC
09H	norAxorBC	A?norBC:xnorBC	C?xorBA:norBA	A?!B:xnorBC	B?xorAC:norAC	A?!C:xnorBC	xnorABC	A?nandBC:xnorBC
0AH	andC!A	A?norBC:C	andCnandBA	A?!B:C	C?!A:andBA	xorCA	xorCandBA	A?nandBC:C
0BH	C?!A:norBA	C?!A:!B	C?nandBA:norBA	C?nandBA:!B	B?xorAC:!A	B?xorAC:nandAC	C?nandBA:xnorBA	nandBxnorAC
0CH	andB!A	A?norBC:B	B?!A:andAC	xorBA	andBnandAC	A?!C:B	xorBandAC	A?nandBC:B
0DH	B?!A:norAC	B?!A:!C	B?!A:xnorAC	C?xorBA:nandBA	B?nandAC:norAC	B?nandAC:!C	B?nandAC:xnorAC	nandCxnorBA
0EH	norAnorBC	xorAorBC	B?!A:C	A?!B:orBC	C?!A:B	A?!C:orBC	B?nandAC:C	A?nandBC:orBC
0FH	!A	nandAorBC	C?nandBA:!A	nandBA	B?nandAC:!A	nandCA	nandAxnorBC	nandABC

Table 5-2 shows the half of 256-entry map corresponding to column index values 8:15.

Table 5-2. Low 8 columns of the 16x16 Map of VPTERNLOG Boolean Logic Operations

Imm	[7:4]							
[3:0]	08H	09H	0AH	0BH	0CH	0DH	0EH	0FH
00H	<i>andABC</i>	<i>andAxnorBC</i>	<i>andCA</i>	<i>B?andAC:A</i>	<i>andBA</i>	<i>C?andBA:A</i>	<i>andAorBC</i>	<i>A</i>
01H	<i>A?andBC:norBC</i>	<i>B?andAC:!C</i>	<i>A?C:norBC</i>	<i>C?A:!B</i>	<i>A?B:norBC</i>	<i>B?A:!C</i>	<i>xnorAorBC</i>	<i>orAnorBC</i>
02H	<i>andCxnorBA</i>	<i>B?andAC:xorAC</i>	<i>B?andAC:C</i>	<i>B?andAC:orAC</i>	<i>C?xnorBA:andBA</i>	<i>B?A:xorAC</i>	<i>B?A:C</i>	<i>B?A:orAC</i>
03H	<i>A?andBC:!B</i>	<i>xnorBandAC</i>	<i>A?C:!B</i>	<i>nandBnandAC</i>	<i>xnorBA</i>	<i>B?A:nandAC</i>	<i>A?orBC:!B</i>	<i>orA!B</i>
04H	<i>andBxnorAC</i>	<i>C?andBA:xorBA</i>	<i>B?xnorAC:andAC</i>	<i>B?xnorAC:A</i>	<i>C?andBA:B</i>	<i>C?andBA:orBA</i>	<i>C?A:B</i>	<i>C?A:orBA</i>
05H	<i>A?andBC:!C</i>	<i>xnorCandBA</i>	<i>xnorCA</i>	<i>C?A:nandBA</i>	<i>A?B:!C</i>	<i>nandCnandBA</i>	<i>A?orBC:!C</i>	<i>orA!C</i>
06H	<i>A?andBC:xorBC</i>	<i>xorABC</i>	<i>A?C:xorBC</i>	<i>B?xnorAC:orAC</i>	<i>A?B:xorBC</i>	<i>C?xnorBA:orBA</i>	<i>A?orBC:xorBC</i>	<i>orAxorBC</i>
07H	<i>xnorAandBC</i>	<i>A?xnorBC:nandBC</i>	<i>A?C:nandBC</i>	<i>nandBxorAC</i>	<i>A?B:nandBC</i>	<i>nandCxorBA</i>	<i>A?orBCnandBC</i>	<i>orAnandBC</i>
08H	<i>andCB</i>	<i>A?xnorBC:andBC</i>	<i>andCorAB</i>	<i>B?C:A</i>	<i>andBorAC</i>	<i>C?B:A</i>	<i>majorABC</i>	<i>orAandBC</i>
09H	<i>B?C:norAC</i>	<i>xnorCB</i>	<i>xnorCorBA</i>	<i>C?orBA:!B</i>	<i>xnorBorAC</i>	<i>B?orAC:!C</i>	<i>A?orBC:xnorBC</i>	<i>orAxnorBC</i>
0AH	<i>A?andBC:C</i>	<i>A?xnorBC:C</i>	<i>C</i>	<i>B?C:orAC</i>	<i>A?B:C</i>	<i>B?orAC:xorAC</i>	<i>orCandBA</i>	<i>orCA</i>
0BH	<i>B?C:!A</i>	<i>B?C:nandAC</i>	<i>orCnorBA</i>	<i>orC!B</i>	<i>B?orAC:!A</i>	<i>B?orAC:nandAC</i>	<i>orCxnorBA</i>	<i>nandBnorAC</i>
0CH	<i>A?andBC:B</i>	<i>A?xnorBC:B</i>	<i>A?C:B</i>	<i>C?orBA:xorBA</i>	<i>B</i>	<i>C?B:orBA</i>	<i>orBandAC</i>	<i>orBA</i>
0DH	<i>C?B!A</i>	<i>C?B:nandBA</i>	<i>C?orBA:!A</i>	<i>C?orBA:nandBA</i>	<i>orBnorAC</i>	<i>orB!C</i>	<i>orBxnorAC</i>	<i>nandCnorBA</i>
0EH	<i>A?andBC:orBC</i>	<i>A?xnorBC:orBC</i>	<i>A?C:orBC</i>	<i>orCxorBA</i>	<i>A?B:orBC</i>	<i>orBxorAC</i>	<i>orCB</i>	<i>orABC</i>
0FH	<i>nandAnandBC</i>	<i>nandAxorBC</i>	<i>orCIA</i>	<i>orCnandBA</i>	<i>orB!A</i>	<i>orBnandAC</i>	<i>nandAnorBC</i>	<i>TRUE</i>

Table 5-1 and Table 5-2 translate each of the possible value of the imm8 byte to a Boolean expression. These tables can also be used by software to translate Boolean expressions to numerical constants to form the imm8 value needed to construct the VPTERNLOG syntax. There is a unique set of three byte constants (FOH, CCH, AAH) that can be used for this purpose as input operands in conjunction with the Boolean expressions defined in those tables. The reverse mapping can be expressed as:

Result_imm8 = Table_Lookup_Entry(FOH, CCH, AAH).

Table_Lookup_Entry is the Boolean expression defined in Table 5-1 and Table 5-2.

5.4 INSTRUCTION SET REFERENCE

<Only instructions modified by AVX512F are included.>

VPERMI2B—Full Permute of Bytes from Two Tables Overwriting the Index

Opcode/ Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEX.DDS.128.66.0F38.W0 75 /r VPERMI2B xmm1 {k1}{z}, xmm2, xmm3/m128	FVM	V/V	AVX512VL AVX512VBMI	Permute bytes in xmm3/m128 and xmm2 using byte indexes in xmm1 and store the byte results in xmm1 using writemask k1.
EVEX.DDS.256.66.0F38.W0 75 /r VPERMI2B ymm1 {k1}{z}, ymm2, ymm3/m256	FVM	V/V	AVX512VL AVX512VBMI	Permute bytes in ymm3/m256 and ymm2 using byte indexes in ymm1 and store the byte results in ymm1 using writemask k1.
EVEX.DDS.512.66.0F38.W0 75 /r VPERMI2B zmm1 {k1}{z}, zmm2, zmm3/m512	FVM	V/V	AVX512VBMI	Permute bytes in zmm3/m512 and zmm2 using byte indexes in zmm1 and store the byte results in zmm1 using writemask k1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
FVM	ModRM:reg (r, w)	EVEX.vvvv	ModRM:r/m (r)	NA

Description

Permutes byte values in the second operand (the first source operand) and the third operand (the second source operand) using the byte indices in the first operand (the destination operand) to select byte elements from the second or third source operands. The selected byte elements are written to the destination at byte granularity under the writemask k1.

The first and second operands are ZMM/YMM/XMM registers. The first operand contains input indices to select elements from the two input tables in the 2nd and 3rd operands. The first operand is also the destination of the result. The third operand can be a ZMM/YMM/XMM register, or a 512/256/128-bit memory location. In each index byte, the id bit for table selection is bit 6/5/4, and bits [5:0]/[4:0]/[3:0] selects element within each input table.

Note that these instructions permit a byte value in the source operands to be copied to more than one location in the destination operand. Also, the same tables can be reused in subsequent iterations, but the index elements are overwritten.

Bits (MAX_VL-1:256/128) of the destination are zeroed for VL=256,128.

Operation**VPERMI2B (EVEX encoded versions)**

(KL, VL) = (16, 128), (32, 256), (64, 512)

IF VL = 128:

id ← 3;

ELSE IF VL = 256:

id ← 4;

ELSE IF VL = 512:

id ← 5;

FI;

TMP_DEST[VL-1:0] ← DEST[VL-1:0];

FOR j ← 0 TO KL-1

off ← 8*SRC1[j*8 + id:j*8];

IF k1[j] OR *no writemask*:

DEST[j*8 + 7:j*8] ← TMP_DEST[j*8+id+1]: SRC2[off+7:off] : SRC1[off+7:off];

ELSE IF *zeroing-masking*

DEST[j*8 + 7:j*8] ← 0;

*ELSE

DEST[j*8 + 7:j*8] remains unchanged*

FI;

ENDFOR

DEST[MAX_VL-1:VL] ← 0;

Intel C/C++ Compiler Intrinsic Equivalent

VPERMI2B __m512i __mm512_permutex2var_epi8(__m512i a, __m512i idx, __m512i b);

VPERMI2B __m512i __mm512_mask2_permutex2var_epi8(__m512i a, __m512i idx, __mmask64 k, __m512i b);

VPERMI2B __m512i __mm512_maskz_permutex2var_epi8(__mmask64 k, __m512i a, __m512i idx, __m512i b);

VPERMI2B __m256i __mm256_permutex2var_epi8(__m256i a, __m256i idx, __m256i b);

VPERMI2B __m256i __mm256_mask2_permutex2var_epi8(__m256i a, __m256i idx, __mmask32 k, __m256i b);

VPERMI2B __m256i __mm256_maskz_permutex2var_epi8(__mmask32 k, __m256i a, __m256i idx, __m256i b);

VPERMI2B __m128i __mm_permutex2var_epi8(__m128i a, __m128i idx, __m128i b);

VPERMI2B __m128i __mm_mask2_permutex2var_epi8(__m128i a, __m128i idx, __mmask16 k, __m128i b);

VPERMI2B __m128i __mm_maskz_permutex2var_epi8(__mmask16 k, __m128i a, __m128i idx, __m128i b);

SIMD Floating-Point Exceptions

None.

Other Exceptions

See Exceptions Type E4NF.nb.

VPERMT2B—Full Permute of Bytes from Two Tables Overwriting a Table

Opcode/ Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEEX.DDS.128.66.0F38.W0 7D /r VPERMT2B xmm1 {k1}{z}, xmm2, xmm3/m128	FVM	V/V	AVX512VL AVX512VBMI	Permute bytes in xmm3/m128 and xmm1 using byte indexes in xmm2 and store the byte results in xmm1 using writemask k1.
EVEEX.NDS.256.66.0F38.W0 7D /r VPERMT2B ymm1 {k1}{z}, ymm2, ymm3/m256	FVM	V/V	AVX512VL AVX512VBMI	Permute bytes in ymm3/m256 and ymm1 using byte indexes in ymm2 and store the byte results in ymm1 using writemask k1.
EVEEX.NDS.512.66.0F38.W0 7D /r VPERMT2B zmm1 {k1}{z}, zmm2, zmm3/m512	FVM	V/V	AVX512VBMI	Permute bytes in zmm3/m512 and zmm1 using byte indexes in zmm2 and store the byte results in zmm1 using writemask k1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
FVM	ModRM:reg (r, w)	EVEX.vvvv	ModRM:r/m (r)	NA

Description

Permutes byte values from two tables, comprising of the first operand (also the destination operand) and the third operand (the second source operand). The second operand (the first source operand) provides byte indices to select byte results from the two tables. The selected byte elements are written to the destination at byte granularity under the writemask k1.

The first and second operands are ZMM/YMM/XMM registers. The second operand contains input indices to select elements from the two input tables in the 1st and 3rd operands. The first operand is also the destination of the result. The second source operand can be a ZMM/YMM/XMM register, or a 512/256/128-bit memory location. In each index byte, the id bit for table selection is bit 6/5/4, and bits [5:0]/[4:0]/[3:0] selects element within each input table.

Note that these instructions permit a byte value in the source operands to be copied to more than one location in the destination operand. Also, the second table and the indices can be reused in subsequent iterations, but the first table is overwritten.

Bits (MAX_VL-1:256/128) of the destination are zeroed for VL=256,128.

Operation

VPERMT2B (EVEX encoded versions)

(KL, VL) = (16, 128), (32, 256), (64, 512)

IF VL = 128:

id ← 3;

ELSE IF VL = 256:

id ← 4;

ELSE IF VL = 512:

id ← 5;

FI;

TMP_DEST[VL-1:0] ← DEST[VL-1:0];

FOR j ← 0 TO KL-1

off ← 8*SRC1[j*8 + id:j*8];

IF k1[j] OR *no writemask*:

DEST[j*8 + 7:j*8] ← SRC1[j*8+id+1]? SRC2[off+7:off] : TMP_DEST[off+7:off];

ELSE IF *zeroing-masking*

DEST[j*8 + 7:j*8] ← 0;

*ELSE

DEST[j*8 + 7:j*8] remains unchanged*

```

    FI;
ENDFOR
DEST[MAX_VL-1:VL] ← 0;

```

Intel C/C++ Compiler Intrinsic Equivalent

```

VPERMT2B __m512i _mm512_permutex2var_epi8(__m512i a, __m512i idx, __m512i b);
VPERMT2B __m512i _mm512_mask_permutex2var_epi8(__m512i a, __mmask64 k, __m512i idx, __m512i b);
VPERMT2B __m512i _mm512_maskz_permutex2var_epi8(__mmask64 k, __m512i a, __m512i idx, __m512i b);
VPERMT2B __m256i _mm256_permutex2var_epi8(__m256i a, __m256i idx, __m256i b);
VPERMT2B __m256i _mm256_mask_permutex2var_epi8(__m256i a, __mmask32 k, __m256i idx, __m256i b);
VPERMT2B __m256i _mm256_maskz_permutex2var_epi8(__mmask32 k, __m256i a, __m256i idx, __m256i b);
VPERMT2B __m128i _mm_permutex2var_epi8(__m128i a, __m128i idx, __m128i b);
VPERMT2B __m128i _mm_mask_permutex2var_epi8(__m128i a, __mmask16 k, __m128i idx, __m128i b);
VPERMT2B __m128i _mm_maskz_permutex2var_epi8(__mmask16 k, __m128i a, __m128i idx, __m128i b);

```

SIMD Floating-Point Exceptions

None.

Other Exceptions

See Exceptions Type E4NF.nb.

VPERMT2W/D/Q/PS/PD—Full Permute from Two Tables Overwriting one Table

Opcode/ Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEX.DDS.128.66.0F38.W1 7D /r VPERMT2W xmm1 {k1}{z}, xmm2, xmm3/m128	FVM	V/V	AVX512VL AVX512BW	Permute word integers from two tables in xmm3/m128 and xmm1 using indexes in xmm2 and store the result in xmm1 using writemask k1.
EVEX.DDS.256.66.0F38.W1 7D /r VPERMT2W ymm1 {k1}{z}, ymm2, ymm3/m256	FVM	V/V	AVX512VL AVX512BW	Permute word integers from two tables in ymm3/m256 and ymm1 using indexes in ymm2 and store the result in ymm1 using writemask k1.
EVEX.DDS.512.66.0F38.W1 7D /r VPERMT2W zmm1 {k1}{z}, zmm2, zmm3/m512	FVM	V/V	AVX512BW	Permute word integers from two tables in zmm3/m512 and zmm1 using indexes in zmm2 and store the result in zmm1 using writemask k1.
EVEX.DDS.128.66.0F38.W0 7E /r VPERMT2D xmm1 {k1}{z}, xmm2, xmm3/m128/m32bcst	FV	V/V	AVX512VL AVX512F	Permute double-words from two tables in xmm3/m128/m32bcst and xmm1 using indexes in xmm2 and store the result in xmm1 using writemask k1.
EVEX.DDS.256.66.0F38.W0 7E /r VPERMT2D ymm1 {k1}{z}, ymm2, ymm3/m256/m32bcst	FV	V/V	AVX512VL AVX512F	Permute double-words from two tables in ymm3/m256/m32bcst and ymm1 using indexes in ymm2 and store the result in ymm1 using writemask k1.
EVEX.DDS.512.66.0F38.W0 7E /r VPERMT2D zmm1 {k1}{z}, zmm2, zmm3/m512/m32bcst	FV	V/V	AVX512F	Permute double-words from two tables in zmm3/m512/m32bcst and zmm1 using indices in zmm2 and store the result in zmm1 using writemask k1.
EVEX.DDS.128.66.0F38.W1 7E /r VPERMT2Q xmm1 {k1}{z}, xmm2, xmm3/m128/m64bcst	FV	V/V	AVX512VL AVX512F	Permute quad-words from two tables in xmm3/m128/m64bcst and xmm1 using indexes in xmm2 and store the result in xmm1 using writemask k1.
EVEX.DDS.256.66.0F38.W1 7E /r VPERMT2Q ymm1 {k1}{z}, ymm2, ymm3/m256/m64bcst	FV	V/V	AVX512VL AVX512F	Permute quad-words from two tables in ymm3/m256/m64bcst and ymm1 using indexes in ymm2 and store the result in ymm1 using writemask k1.
EVEX.DDS.512.66.0F38.W1 7E /r VPERMT2Q zmm1 {k1}{z}, zmm2, zmm3/m512/m64bcst	FV	V/V	AVX512F	Permute quad-words from two tables in zmm3/m512/m64bcst and zmm1 using indices in zmm2 and store the result in zmm1 using writemask k1.
EVEX.DDS.128.66.0F38.W0 7F /r VPERMT2PS xmm1 {k1}{z}, xmm2, xmm3/m128/m32bcst	FV	V/V	AVX512VL AVX512F	Permute single-precision FP values from two tables in xmm3/m128/m32bcst and xmm1 using indexes in xmm2 and store the result in xmm1 using writemask k1.
EVEX.DDS.256.66.0F38.W0 7F /r VPERMT2PS ymm1 {k1}{z}, ymm2, ymm3/m256/m32bcst	FV	V/V	AVX512VL AVX512F	Permute single-precision FP values from two tables in ymm3/m256/m32bcst and ymm1 using indexes in ymm2 and store the result in ymm1 using writemask k1.
EVEX.DDS.512.66.0F38.W0 7F /r VPERMT2PS zmm1 {k1}{z}, zmm2, zmm3/m512/m32bcst	FV	V/V	AVX512F	Permute single-precision FP values from two tables in zmm3/m512/m32bcst and zmm1 using indices in zmm2 and store the result in zmm1 using writemask k1.
EVEX.DDS.128.66.0F38.W1 7F /r VPERMT2PD xmm1 {k1}{z}, xmm2, xmm3/m128/m64bcst	FV	V/V	AVX512VL AVX512F	Permute double-precision FP values from two tables in xmm3/m128/m64bcst and xmm1 using indexes in xmm2 and store the result in xmm1 using writemask k1.
EVEX.DDS.256.66.0F38.W1 7F /r VPERMT2PD ymm1 {k1}{z}, ymm2, ymm3/m256/m64bcst	FV	V/V	AVX512VL AVX512F	Permute double-precision FP values from two tables in ymm3/m256/m64bcst and ymm1 using indexes in ymm2 and store the result in ymm1 using writemask k1.
EVEX.DDS.512.66.0F38.W1 7F /r VPERMT2PD zmm1 {k1}{z}, zmm2, zmm3/m512/m64bcst	FV	V/V	AVX512F	Permute double-precision FP values from two tables in zmm3/m512/m64bcst and zmm1 using indices in zmm2 and store the result in zmm1 using writemask k1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
FVM	ModRM:reg (r,w)	EVEX.vvvv (r)	ModRM:r/m (r)	NA
FV	ModRM:reg (r, w)	EVEX.vvvv (r)	ModRM:r/m (r)	NA

Description

Permutes 16-bit/32-bit/64-bit values in the first operand and the third operand (the second source operand) using indices in the second operand (the first source operand) to select elements from the first and third operands. The selected elements are written to the destination operand (the first operand) according to the writemask k1.

The first and second operands are ZMM/YMM/XMM registers. The second operand contains input indices to select elements from the two input tables in the 1st and 3rd operands. The first operand is also the destination of the result.

D/Q/PS/PD element versions: The second source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location or a 512/256/128-bit vector broadcasted from a 32/64-bit memory location. Broadcast from the low 32/64-bit memory location is performed if EVEX.b and the id bit for table selection are set (selecting table_2).

Dword/PS versions: The id bit for table selection is bit 4/3/2, depending on VL=512, 256, 128. Bits [3:0]/[2:0]/[1:0] of each element in the input index vector select an element within the two source operands. If the id bit is 0, table_1 (the first source) is selected; otherwise the second source operand is selected.

Qword/PD versions: The id bit for table selection is bit 3/2/1, and bits [2:0]/[1:0] /bit 0 selects element within each input table.

Word element versions: The second source operand can be a ZMM/YMM/XMM register, or a 512/256/128-bit memory location. The id bit for table selection is bit 5/4/3, and bits [4:0]/[3:0]/[2:0] selects element within each input table.

Note that these instructions permit a 16-bit/32-bit/64-bit value in the source operands to be copied to more than one location in the destination operand. Note also that in this case, the same index can be reused for example for a second iteration, while the table elements being permuted are overwritten.

Bits (MAX_VL-1:256/128) of the destination are zeroed for VL=256,128.

Operation

VPERMT2W (EVEX encoded versions)

(KL, VL) = (8, 128), (16, 256), (32, 512)

IF VL = 128

id ← 2

FI;

IF VL = 256

id ← 3

FI;

IF VL = 512

id ← 4

FI;

TMP_DEST ← DEST

FOR j ← 0 TO KL-1

i ← j * 16

off ← 16*SRC1[j+id:i]

IF k1[j] OR *no writemask*

THEN

DEST[i+15:i]=SRC1[j+id+1] ? SRC2[off+15:off]
: TMP_DEST[off+15:off]

ELSE

IF *merging-masking* ; merging-masking

THEN *DEST[i+15:i] remains unchanged*

ELSE ; zeroing-masking

```

        DEST[i+15:i] ← 0
    FI
FI;
ENDFOR
DEST[MAX_VL-1:VL] ← 0

VPERMT2D/VPERMT2PS (EVEX encoded versions)
(KL, VL) = (4, 128), (8, 256), (16, 512)
IF VL = 128
    id ← 1
FI;
IF VL = 256
    id ← 2
FI;
IF VL = 512
    id ← 3
FI;
TMP_DEST ← DEST
FOR j ← 0 TO KL-1
    i ← j * 32
    off ← 32 * SRC1[i+id:i]
    IF k1[j] OR *no writemask*
        THEN
            IF (EVEX.b = 1) AND (SRC2 *is memory*)
                THEN
                    DEST[i+31:i] ← SRC1[i+id+1] ? SRC2[31:0]
                    : TMP_DEST[off+31:off]
                ELSE
                    DEST[i+31:i] ← SRC1[i+id+1] ? SRC2[off+31:off]
                    : TMP_DEST[off+31:off]
                FI
            ELSE
                IF *merging-masking* ; merging-masking
                    THEN *DEST[i+31:i] remains unchanged*
                ELSE ; zeroing-masking
                    DEST[i+31:i] ← 0
                FI
            FI
        FI
    FI;
ENDFOR
DEST[MAX_VL-1:VL] ← 0

```

VPERMT2Q/VPERMT2PD (EVEX encoded versions)

(KL, VL) = (2, 128), (4, 256), (8, 512)

```

IF VL = 128
    id ← 0
FI;
IF VL = 256
    id ← 1
FI;
IF VL = 512
    id ← 2
FI;
TMP_DEST ← DEST
FOR j ← 0 TO KL-1

```

```

i ← j * 64
off ← 64 * SRC1[i+id:i]
IF k1[j] OR *no writemask*
    THEN
        IF (EVEX.b = 1) AND (SRC2 *is memory*)
            THEN
                DEST[i+63:i] ← SRC1[i+id+1] ? SRC2[63:0]
                : TMP_DEST[off+63:off]
            ELSE
                DEST[i+63:i] ← SRC1[i+id+1] ? SRC2[off+63:off]
                : TMP_DEST[off+63:off]
            FI
        ELSE
            IF *merging-masking* ; merging-masking
                THEN *DEST[i+63:i] remains unchanged*
            ELSE ; zeroing-masking
                DEST[i+63:i] ← 0
            FI
        FI
    FI;
ENDFOR
DEST[MAX_VL-1:VL] ← 0

```

Intel C/C++ Compiler Intrinsic Equivalent

```

VPERMT2D __m512i _mm512_permutex2var_epi32(__m512i a, __m512i idx, __m512i b);
VPERMT2D __m512i _mm512_mask_permutex2var_epi32(__m512i a, __mmask16 k, __m512i idx, __m512i b);
VPERMT2D __m512i _mm512_mask2_permutex2var_epi32(__m512i a, __m512i idx, __mmask16 k, __m512i b);
VPERMT2D __m512i _mm512_maskz_permutex2var_epi32(__mmask16 k, __m512i a, __m512i idx, __m512i b);
VPERMT2D __m256i _mm256_permutex2var_epi32(__m256i a, __m256i idx, __m256i b);
VPERMT2D __m256i _mm256_mask_permutex2var_epi32(__m256i a, __mmask8 k, __m256i idx, __m256i b);
VPERMT2D __m256i _mm256_mask2_permutex2var_epi32(__m256i a, __m256i idx, __mmask8 k, __m256i b);
VPERMT2D __m256i _mm256_maskz_permutex2var_epi32(__mmask8 k, __m256i a, __m256i idx, __m256i b);
VPERMT2D __m128i _mm_permutex2var_epi32(__m128i a, __m128i idx, __m128i b);
VPERMT2D __m128i _mm_mask_permutex2var_epi32(__m128i a, __mmask8 k, __m128i idx, __m128i b);
VPERMT2D __m128i _mm_mask2_permutex2var_epi32(__m128i a, __m128i idx, __mmask8 k, __m128i b);
VPERMT2D __m128i _mm_maskz_permutex2var_epi32(__mmask8 k, __m128i a, __m128i idx, __m128i b);
VPERMT2PD __m512d _mm512_permutex2var_pd(__m512d a, __m512i idx, __m512d b);
VPERMT2PD __m512d _mm512_mask_permutex2var_pd(__m512d a, __mmask8 k, __m512i idx, __m512d b);
VPERMT2PD __m512d _mm512_mask2_permutex2var_pd(__m512d a, __m512i idx, __mmask8 k, __m512d b);
VPERMT2PD __m512d _mm512_maskz_permutex2var_pd(__mmask8 k, __m512d a, __m512i idx, __m512d b);
VPERMT2PD __m256d _mm256_permutex2var_pd(__m256d a, __m256i idx, __m256d b);
VPERMT2PD __m256d _mm256_mask_permutex2var_pd(__m256d a, __mmask8 k, __m256i idx, __m256d b);
VPERMT2PD __m256d _mm256_mask2_permutex2var_pd(__m256d a, __m256i idx, __mmask8 k, __m256d b);
VPERMT2PD __m256d _mm256_maskz_permutex2var_pd(__mmask8 k, __m256d a, __m256i idx, __m256d b);
VPERMT2PD __m128d _mm_permutex2var_pd(__m128d a, __m128i idx, __m128d b);
VPERMT2PD __m128d _mm_mask_permutex2var_pd(__m128d a, __mmask8 k, __m128i idx, __m128d b);
VPERMT2PD __m128d _mm_mask2_permutex2var_pd(__m128d a, __m128i idx, __mmask8 k, __m128d b);
VPERMT2PD __m128d _mm_maskz_permutex2var_pd(__mmask8 k, __m128d a, __m128i idx, __m128d b);
VPERMT2PS __m512 _mm512_permutex2var_ps(__m512 a, __m512i idx, __m512 b);
VPERMT2PS __m512 _mm512_mask_permutex2var_ps(__m512 a, __mmask16 k, __m512i idx, __m512 b);
VPERMT2PS __m512 _mm512_mask2_permutex2var_ps(__m512 a, __m512i idx, __mmask16 k, __m512 b);
VPERMT2PS __m512 _mm512_maskz_permutex2var_ps(__mmask16 k, __m512 a, __m512i idx, __m512 b);

```

```

VPERMT2PS __m256 __mm256_permutex2var_ps(__m256 a, __m256i idx, __m256 b);
VPERMT2PS __m256 __mm256_mask_permutex2var_ps(__m256 a, __mmask8 k, __m256i idx, __m256 b);
VPERMT2PS __m256 __mm256_mask2_permutex2var_ps(__m256 a, __m256i idx, __mmask8 k, __m256 b);
VPERMT2PS __m256 __mm256_maskz_permutex2var_ps(__mmask8 k, __m256 a, __m256i idx, __m256 b);
VPERMT2PS __m128 __mm_permutex2var_ps(__m128 a, __m128i idx, __m128 b);
VPERMT2PS __m128 __mm_mask_permutex2var_ps(__m128 a, __mmask8 k, __m128i idx, __m128 b);
VPERMT2PS __m128 __mm_mask2_permutex2var_ps(__m128 a, __m128i idx, __mmask8 k, __m128 b);
VPERMT2PS __m128 __mm_maskz_permutex2var_ps(__mmask8 k, __m128 a, __m128i idx, __m128 b);
VPERMT2Q __m512i __mm512_permutex2var_epi64(__m512i a, __m512i idx, __m512i b);
VPERMT2Q __m512i __mm512_mask_permutex2var_epi64(__m512i a, __mmask8 k, __m512i idx, __m512i b);
VPERMT2Q __m512i __mm512_mask2_permutex2var_epi64(__m512i a, __m512i idx, __mmask8 k, __m512i b);
VPERMT2Q __m512i __mm512_maskz_permutex2var_epi64(__mmask8 k, __m512i a, __m512i idx, __m512i b);
VPERMT2Q __m256i __mm256_permutex2var_epi64(__m256i a, __m256i idx, __m256i b);
VPERMT2Q __m256i __mm256_mask_permutex2var_epi64(__m256i a, __mmask8 k, __m256i idx, __m256i b);
VPERMT2Q __m256i __mm256_mask2_permutex2var_epi64(__m256i a, __m256i idx, __mmask8 k, __m256i b);
VPERMT2Q __m256i __mm256_maskz_permutex2var_epi64(__mmask8 k, __m256i a, __m256i idx, __m256i b);
VPERMT2Q __m128i __mm_permutex2var_epi64(__m128i a, __m128i idx, __m128i b);
VPERMT2Q __m128i __mm_mask_permutex2var_epi64(__m128i a, __mmask8 k, __m128i idx, __m128i b);
VPERMT2Q __m128i __mm_mask2_permutex2var_epi64(__m128i a, __m128i idx, __mmask8 k, __m128i b);
VPERMT2Q __m128i __mm_maskz_permutex2var_epi64(__mmask8 k, __m128i a, __m128i idx, __m128i b);
VPERMT2W __m512i __mm512_permutex2var_epi16(__m512i a, __m512i idx, __m512i b);
VPERMT2W __m512i __mm512_mask_permutex2var_epi16(__m512i a, __mmask32 k, __m512i idx, __m512i b);
VPERMT2W __m512i __mm512_mask2_permutex2var_epi16(__m512i a, __m512i idx, __mmask32 k, __m512i b);
VPERMT2W __m512i __mm512_maskz_permutex2var_epi16(__mmask32 k, __m512i a, __m512i idx, __m512i b);
VPERMT2W __m256i __mm256_permutex2var_epi16(__m256i a, __m256i idx, __m256i b);
VPERMT2W __m256i __mm256_mask_permutex2var_epi16(__m256i a, __mmask16 k, __m256i idx, __m256i b);
VPERMT2W __m256i __mm256_mask2_permutex2var_epi16(__m256i a, __m256i idx, __mmask16 k, __m256i b);
VPERMT2W __m256i __mm256_maskz_permutex2var_epi16(__mmask16 k, __m256i a, __m256i idx, __m256i b);
VPERMT2W __m128i __mm_permutex2var_epi16(__m128i a, __m128i idx, __m128i b);
VPERMT2W __m128i __mm_mask_permutex2var_epi16(__m128i a, __mmask8 k, __m128i idx, __m128i b);
VPERMT2W __m128i __mm_mask2_permutex2var_epi16(__m128i a, __m128i idx, __mmask8 k, __m128i b);
VPERMT2W __m128i __mm_maskz_permutex2var_epi16(__mmask8 k, __m128i a, __m128i idx, __m128i b);

```

SIMD Floating-Point Exceptions

None.

Other Exceptions

VPERMT2D/Q/PS/PD: See Exceptions Type E4NF.

VPERMT2W: See Exceptions Type E4NF.nb.

VPMADD52LUQ—Packed Multiply of Unsigned 52-bit Integers and Add the Low 52-bit Products to Qword Accumulators

Opcode/ Instruction	Op/En	32/64 bit Mode Support	CPUID	Description
EVEX.DDS.128.66.0F38.W1 B4 /r VPMADD52LUQ xmm1 {k1}{z}, xmm2,xmm3/m128/m64bcst	FV	V/V	AVX512IFMA AVX512VL	Multiply unsigned 52-bit integers in xmm2 and xmm3/m128 and add the low 52 bits of the 104-bit product to the qword unsigned integers in xmm1 using writemask k1.
EVEX.DDS.256.66.0F38.W1 B4 /r VPMADD52LUQ ymm1 {k1}{z}, ymm2, ymm3/m256/m64bcst	FV	V/V	AVX512IFMA AVX512VL	Multiply unsigned 52-bit integers in ymm2 and ymm3/m128 and add the low 52 bits of the 104-bit product to the qword unsigned integers in ymm1 using writemask k1.
EVEX.DDS.512.66.0F38.W1 B4 /r VPMADD52LUQ zmm1 {k1}{z}, zmm2,zmm3/m512/m64bcst	FV	V/V	AVX512IFMA	Multiply unsigned 52-bit integers in zmm2 and zmm3/m128 and add the low 52 bits of the 104-bit product to the qword unsigned integers in zmm1 using writemask k1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
FV	ModRM:reg (r, w)	EVEX.vvvv (r)	ModRM:r/m(r)	NA

Description

Multiplies packed unsigned 52-bit integers in each qword element of the first source operand (the second operand) with the packed unsigned 52-bit integers in the corresponding elements of the second source operand (the third operand) to form packed 104-bit intermediate results. The low 52-bit, unsigned integer of each 104-bit product is added to the corresponding qword unsigned integer of the destination operand (the first operand) under the writemask k1.

The first source operand is a ZMM/YMM/XMM register. The second source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location or a 512/256/128-bit vector broadcasted from a 64-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1 at 64-bit granularity.

Operation

VPMADD52LUQ (EVEX encoded)

(KL, VL) = (2, 128), (4, 256), (8, 512)

FOR j ← 0 TO KL-1

 i ← j * 64;

 IF k1[j] OR *no writemask* THEN

 IF src2 is Memory AND EVEX.b=1 THEN

 tsrc2[63:0] ← ZeroExtend64(src2[51:0]);

 ELSE

 tsrc2[63:0] ← ZeroExtend64(src2[i+51:i]);

 FI;

 Temp128[127:0] ← ZeroExtend64(src1[i+51:i]) * tsrc2[63:0];

 Temp2[63:0] ← DEST[i+63:i] + ZeroExtend64(temp128[51:0]);

 DEST[i+63:i] ← Temp2[63:0];

 ELSE

 IF *zeroing-masking* THEN

 DEST[i+63:i] ← 0;

 ELSE *merge-masking*

 DEST[i+63:i] is unchanged;

 FI;

 FI;

ENDFOR

DEST[MAX_VL-1:VL] ← 0;

Intel C/C++ Compiler Intrinsic Equivalent

VPMADD52LUQ __m512i __mm512_madd52lo_epu64(__m512i a, __m512i b, __m512i c);

VPMADD52LUQ __m512i __mm512_mask_madd52lo_epu64(__m512i s, __mmask8 k, __m512i a, __m512i b, __m512i c);

VPMADD52LUQ __m512i __mm512_maskz_madd52lo_epu64(__mmask8 k, __m512i a, __m512i b, __m512i c);

VPMADD52LUQ __m256i __mm256_madd52lo_epu64(__m256i a, __m256i b, __m256i c);

VPMADD52LUQ __m256i __mm256_mask_madd52lo_epu64(__m256i s, __mmask8 k, __m256i a, __m256i b, __m256i c);

VPMADD52LUQ __m256i __mm256_maskz_madd52lo_epu64(__mmask8 k, __m256i a, __m256i b, __m256i c);

VPMADD52LUQ __m128i __mm_madd52lo_epu64(__m128i a, __m128i b, __m128i c);

VPMADD52LUQ __m128i __mm_mask_madd52lo_epu64(__m128i s, __mmask8 k, __m128i a, __m128i b, __m128i c);

VPMADD52LUQ __m128i __mm_maskz_madd52lo_epu64(__mmask8 k, __m128i a, __m128i b, __m128i c);

Flags Affected

None.

SIMD Floating-Point Exceptions

None

Other Exceptions

See Exceptions Type E4.

VPMADD52HUQ—Packed Multiply of Unsigned 52-bit Unsigned Integers and Add High 52-bit Products to 64-bit Accumulators

Opcode/ Instruction	Op/En	32/64 bit Mode Support	CPUID	Description
EVEX.DDS.128.66.0F38.W1 B5 /r VPMADD52HUQ xmm1 {k1}{z}, xmm2, xmm3/m128/m64bcst	FV	V/V	AVX512IFMA AVX512VL	Multiply unsigned 52-bit integers in xmm2 and xmm3/m128 and add the high 52 bits of the 104-bit product to the qword unsigned integers in xmm1 using writemask k1.
EVEX.DDS.256.66.0F38.W1 B5 /r VPMADD52HUQ ymm1 {k1}{z}, ymm2, ymm3/m256/m64bcst	FV	V/V	AVX512IFMA AVX512VL	Multiply unsigned 52-bit integers in ymm2 and ymm3/m128 and add the high 52 bits of the 104-bit product to the qword unsigned integers in ymm1 using writemask k1.
EVEX.DDS.512.66.0F38.W1 B5 /r VPMADD52HUQ zmm1 {k1}{z}, zmm2, zmm3/m512/m64bcst	FV	V/V	AVX512IFMA	Multiply unsigned 52-bit integers in zmm2 and zmm3/m128 and add the high 52 bits of the 104-bit product to the qword unsigned integers in zmm1 using writemask k1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
FV	ModRM:reg (r, w)	EVEX.vvvv (r)	ModRM:r/m(r)	NA

Description

Multiplies packed unsigned 52-bit integers in each qword element of the first source operand (the second operand) with the packed unsigned 52-bit integers in the corresponding elements of the second source operand (the third operand) to form packed 104-bit intermediate results. The high 52-bit, unsigned integer of each 104-bit product is added to the corresponding qword unsigned integer of the destination operand (the first operand) under the writemask k1.

The first source operand is a ZMM/YMM/XMM register. The second source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location or a 512/256/128-bit vector broadcasted from a 64-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1 at 64-bit granularity.

Operation

VPMADD52HUQ (EVEX encoded)

(KL, VL) = (2, 128), (4, 256), (8, 512)

FOR j ← 0 TO KL-1

 i ← j * 64;

 IF k1[j] OR *no writemask* THEN

 IF src2 is Memory AND EVEX.b=1 THEN

 tsrc2[63:0] ← ZeroExtend64(src2[51:0]);

 ELSE

 tsrc2[63:0] ← ZeroExtend64(src2[i+51:i]);

 FI;

 Temp128[127:0] ← ZeroExtend64(src1[i+51:i]) * tsrc2[63:0];

 Temp2[63:0] ← DEST[i+63:i] + ZeroExtend64(temp128[103:52]);

 DEST[i+63:i] ← Temp2[63:0];

 ELSE

 IF *zeroing-masking* THEN

 DEST[i+63:i] ← 0;

 ELSE *merge-masking*

 DEST[i+63:i] is unchanged;

 FI;

 FI;

ENDFOR

DEST[MAX_VL-1:VL] ← 0

Intel C/C++ Compiler Intrinsic Equivalent

VPMADD52HUQ __m512i __mm512_madd52hi_epu64(__m512i a, __m512i b, __m512i c);

VPMADD52HUQ __m512i __mm512_mask_madd52hi_epu64(__m512i s, __mmask8 k, __m512i a, __m512i b, __m512i c);

VPMADD52HUQ __m512i __mm512_maskz_madd52hi_epu64(__mmask8 k, __m512i a, __m512i b, __m512i c);

VPMADD52HUQ __m256i __mm256_madd52hi_epu64(__m256i a, __m256i b, __m256i c);

VPMADD52HUQ __m256i __mm256_mask_madd52hi_epu64(__m256i s, __mmask8 k, __m256i a, __m256i b, __m256i c);

VPMADD52HUQ __m256i __mm256_maskz_madd52hi_epu64(__mmask8 k, __m256i a, __m256i b, __m256i c);

VPMADD52HUQ __m128i __mm_madd52hi_epu64(__m128i a, __m128i b, __m128i c);

VPMADD52HUQ __m128i __mm_mask_madd52hi_epu64(__m128i s, __mmask8 k, __m128i a, __m128i b, __m128i c);

VPMADD52HUQ __m128i __mm_maskz_madd52hi_epu64(__mmask8 k, __m128i a, __m128i b, __m128i c);

Flags Affected

None.

SIMD Floating-Point Exceptions

None

Other Exceptions

See Exceptions Type E4.

VPMULTISHIFTQB – Select Packed Unaligned Bytes from Quadword Sources

Opcode / Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEX.NDS.128.66.0F38.W1 83 /r VPMULTISHIFTQB xmm1 {k1}{z}, xmm2,xmm3/m128/m64bcst	FV	V/V	AVX512VBMI AVX512VL	Select unaligned bytes from qwords in xmm3/m128/m64bcst using control bytes in xmm2, write byte results to xmm1 under k1.
EVEX.NDS.256.66.0F38.W1 83 /r VPMULTISHIFTQB ymm1 {k1}{z}, ymm2,ymm3/m256/m64bcst	FV	V/V	AVX512VBMI AVX512VL	Select unaligned bytes from qwords in ymm3/m256/m64bcst using control bytes in ymm2, write byte results to ymm1 under k1.
EVEX.NDS.512.66.0F38.W1 83 /r VPMULTISHIFTQB zmm1 {k1}{z}, zmm2,zmm3/m512/m64bcst	FV	V/V	AVX512VBMI	Select unaligned bytes from qwords in zmm3/m512/m64bcst using control bytes in zmm2, write byte results to zmm1 under k1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
FV	ModRM:reg (w)	EVEX.vvvv (r)	ModRM:r/m (r)	NA

Description

This instruction selects eight unaligned bytes from each input qword element of the second source operand (the third operand) and writes eight assembled bytes for each qword element in the destination operand (the first operand). Each byte result is selected using a byte-granular shift control within the corresponding qword element of the first source operand (the second operand). Each byte result in the destination operand is updated under the writemask k1.

Only the low 6 bits of each control byte are used to select an 8-bit slot to extract the output byte from the qword data in the second source operand. The starting bit of the 8-bit slot can be unaligned relative to any byte boundary and is left-shifted from the beginning of the input qword source by the amount specified in the low 6-bit of the control byte. If the 8-bit slot would exceed the qword boundary, the out-of-bound portion of the 8-bit slot is wrapped back to start from bit 0 of the input qword element.

The first source operand is a ZMM/YMM/XMM register. The second source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location or a 512/256/128-bit vector broadcasted from a 64-bit memory location. The destination operand is a ZMM/YMM/XMM register.

Operation

VPMULTISHIFTQB DEST, SRC1, SRC2 (EVEX encoded version)

(KL, VL) = (2, 128), (4, 256), (8, 512)

```

FOR i ← 0 TO KL-1
    q ← i * 64;
    IF src2 is Memory AND EVEX.b=1 THEN
        tcur64[63:0] ← src2[63:0];
    ELSE
        tcur64[63:0] ← src2[q+63:q];
    FI;
    FOR j ← 0 to 7 // iterate each byte in qword
        ctrl ← src1[q+j*8+7:q+j*8] & 63;
        FOR k ← 0 to 7 // iterate each bit in byte
            tmp8[k] ← tcur64[(ctrl+k) & 63];
        ENDFOR
        IF k1[i*8+j] or no writemask THEN
            dst[q+ j*8 + 7: q + j*8] ← tmp8[7:0];
        ELSE IF zeroing-masking THEN
            dst[q+ j*8 + 7: q + j*8] ← 0;
        ENDFOR
    ENDFOR
DEST[MAX_VL-1:VL] ← 0;

```

Intel C/C++ Compiler Intrinsic Equivalent

```

VPMULTISHIFTQB __m512i __mm512_multishift_epi64_epi8( __m512i a, __m512i b);
VPMULTISHIFTQB __m512i __mm512_mask_multishift_epi64_epi8(__m512i s, __mmask64 k, __m512i a, __m512i b);
VPMULTISHIFTQB __m512i __mm512_maskz_multishift_epi64_epi8( __mmask64 k, __m512i a, __m512i b);
VPMULTISHIFTQB __m256i __mm256_multishift_epi64_epi8( __m256i a, __m256i b);
VPMULTISHIFTQB __m256i __mm256_mask_multishift_epi64_epi8(__m256i s, __mmask32 k, __m256i a, __m256i b);
VPMULTISHIFTQB __m256i __mm256_maskz_multishift_epi64_epi8( __mmask32 k, __m256i a, __m256i b);
VPMULTISHIFTQB __m128i __mm_multishift_epi64_epi8( __m128i a, __m128i b);
VPMULTISHIFTQB __m128i __mm_mask_multishift_epi64_epi8(__m128i s, __mmask8 k, __m128i a, __m128i b);
VPMULTISHIFTQB __m128i __mm_maskz_multishift_epi64_epi8( __mmask8 k, __m128i a, __m128i b);

```

SIMD Floating-Point Exceptions

None.

Other Exceptions

See Exceptions Type E4NF.

CHAPTER 6

ADDITIONAL 512-BIT INSTRUCTION EXTENSIONS

This chapter describes additional 512-bit instruction extensions. Instructions use the same notations and conventions listed in *Section 5.1* and *Section 5.1.5.1*.

6.1 INSTRUCTION SET REFERENCE

V4FMADDPS/V4FNMADDPS — Packed Single-Precision Floating-Point Fused Multiply-Add (4-iterations)

Opcode/ Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEX.DDS.512.F2.0F38.W0 9A /r V4FMADDPS zmm1{k1}{z}, zmm2+3, m128	T1_4X	V/V	AVX512_4FMAPS	Multiply packed single-precision floating-point values from source register block indicated by zmm2 by values from m128 and accumulate the result in zmm1.
EVEX.DDS.512.F2.0F38.W0 AA /r V4FNMADDPS zmm1{k1}{z}, zmm2+3, m128	T1_4X	V/V	AVX512_4FMAPS	Multiply and negate packed single-precision floating-point values from source register block indicated by zmm2 by values from m128 and accumulate the result in zmm1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
T1_4X	ModRM:reg (r, w)	EVEX.vvvv	ModRM:r/m (r)	NA

Description

This instruction computes 4 sequential packed fused single-precision floating-point multiply-add instructions with a sequentially selected memory operand in each of the four steps.

In the above box, the notation of “+3” is used to denote that the instruction accesses 4 source registers based on that operand; sources are consecutive, start in a multiple of 4 boundary, and contain the encoded register operand.

This instruction supports memory fault suppression. The entire memory operand is loaded if any of the 16 lowest significant mask bits is set to 1 or if a “no masking” encoding is used.

The tuple type T1_4X implies that 4 32-bit elements (16 bytes) are referenced by the memory operation portion of this instruction.

Rounding is performed at every FMA (fused multiply and add) boundary. Exceptions are also taken sequentially. Pre- and post-computational exceptions of the first FMA take priority over the pre- and post-computational exceptions of the second FMA, etc.

Operation

src_reg_id is the 5 bit index of the vector register specified in the instruction as the src1 register.

define NFMA_PS(kl, vl, dest, k1, msrc, regs_loaded, src_base, posneg):

tmpdest ← dest

// reg[] is an array representing the SIMD register file.

for j ← 0 to regs_loaded-1:

for i ← 0 to kl-1:

if k1[i] or *no writemask*:

if posneg = 0:

tmpdest.single[i] ← RoundFPControl_MXCSR(tmpdest.single[i] - reg[src_base + j].single[i] * msrc.single[j])

else:

tmpdest.single[i] ← RoundFPControl_MXCSR(tmpdest.single[i] + reg[src_base + j].single[i] * msrc.single[j])

else if *zeroing*:

tmpdest.single[i] ← 0

dest ← tmpdst

dest[MAX_VL-1:VL] ← 0

V4FMADDPS and V4FNMADDPS dest{k1}, src1, msrc (AVX512)

kl, vl = (16, 512)

regs_loaded ← 4

src_base ← src_reg_id & ~3 // for src1 operand

posneg ← 0 if negative form, 1 otherwise

NFMA_PS(kl, vl, dest, k1, msrc, regs_loaded, src_base, posneg)

Intel C/C++ Compiler Intrinsic Equivalent

V4FMADDPS __m512 __mm512_4fmadd_ps(__m512, __m512x4, __m128 *);

V4FMADDPS __m512 __mm512_mask_4fmadd_ps(__m512, __mmask16, __m512x4, __m128 *);

V4FMADDPS __m512 __mm512_maskz_4fmadd_ps(__mmask16, __m512, __m512x4, __m128 *);

V4FNMADDPS __m512 __mm512_4fnmadd_ps(__m512, __m512x4, __m128 *);

V4FNMADDPS __m512 __mm512_mask_4fnmadd_ps(__m512, __mmask16, __m512x4, __m128 *);

V4FNMADDPS __m512 __mm512_maskz_4fnmadd_ps(__mmask16, __m512, __m512x4, __m128 *);

SIMD Floating-Point Exceptions

Overflow, Underflow, Invalid, Precision, Denormal.

Other Exceptions

See Type E2; additionally

#UD If the EVEX broadcast bit is set to 1.

#UD If the MODRM.mod = 0b11.

V4FMADDSS/V4FNMADDSS –Scalar Single-Precision Floating-Point Fused Multiply-Add (4-iterations)

Opcode/ Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEX.DDS.LLIG.F2.0F38.W0 9B /r V4FMADDSS xmm1{k1}{z}, xmm2+3, m128	T1_4X	V/V	AVX512_4FMAPS	Multiply scalar single-precision floating-point values from source register block indicated by xmm2 by values from m128 and accumulate the result in xmm1.
EVEX.DDS.LLIG.F2.0F38.W0 AB /r V4FNMADDSS xmm1{k1}{z}, xmm2+3, m128	T1_4X	V/V	AVX512_4FMAPS	Multiply and negate scalar single-precision floating-point values from source register block indicated by xmm2 by values from m128 and accumulate the result in xmm1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
T1_4X	ModRM:reg (r, w)	EVEX.vvvv	ModRM:r/m (r)	NA

Description

This instruction computes 4 sequential scalar fused single-precision floating-point multiply-add instructions with a sequentially selected memory operand in each of the four steps.

In the above box, the notation of “+3” is used to denote that the instruction accesses 4 source registers based that operand; sources are consecutive, start in a multiple of 4 boundary, and contain the encoded register operand.

This instruction supports memory fault suppression. The entire memory operand is loaded if the least significant mask bit is set to 1 or if a “no masking” encoding is used.

The tuple type T1_4X implies that 4 32-bit elements (16 bytes) are referenced by the memory operation portion of this instruction.

Rounding is performed at every FMA boundary. Exceptions are also taken sequentially. Pre- and post-computational exceptions of the first FMA take priority over the pre- and post-computational exceptions of the second FMA, etc.

Operation

src_reg_id is the 5 bit index of the vector register specified in the instruction as the src1 register.

```

define NFMA_SS(vl, dest, k1, msrc, regs_loaded, src_base, posneg):
    tmpdest ← dest
    // reg[] is an array representing the SIMD register file.
    if k1[0] or *no writemask*:
        for j ← 0 to regs_loaded - 1:
            if posneg = 0:
                tmpdest.single[0] ← RoundFPControl_MXCSR(tmpdest.single[0] - reg[src_base + j].single[0] * msrc.single[j])
            else:
                tmpdest.single[0] ← RoundFPControl_MXCSR(tmpdest.single[0] + reg[src_base + j].single[0] * msrc.single[j])
    else if *zeroing*:
        tmpdest.single[0] ← 0
    dest ← tmpdst
    dest[MAX_VL-1:VL] ← 0

```

ADDITIONAL 512-BIT INSTRUCTION EXTENSIONS

V4FMADDSS and V4FNMADDSS dest{k1}, src1, msrc (AVX512)
vl = 128

regs_loaded \leftarrow 4
src_base \leftarrow src_reg_id & ~3 // for src1 operand
posneg \leftarrow 0 if negative form, 1 otherwise
NFMA_SS(vl, dest, k1, msrc, regs_loaded, src_base, posneg)

Intel C/C++ Compiler Intrinsic Equivalent

V4FMADDSS __m128 _mm_4fmadd_ss(__m128, __m128x4, __m128 *);
V4FMADDSS __m128 _mm_mask_4fmadd_ss(__m128, __mmask8, __m128x4, __m128 *);
V4FMADDSS __m128 _mm_maskz_4fmadd_ss(__mmask8, __m128, __m128x4, __m128 *);
V4FNMADDSS __m128 _mm_4fnmadd_ss(__m128, __m128x4, __m128 *);
V4FNMADDSS __m128 _mm_mask_4fnmadd_ss(__m128, __mmask8, __m128x4, __m128 *);
V4FNMADDSS __m128 _mm_maskz_4fnmadd_ss(__mmask8, __m128, __m128x4, __m128 *);

SIMD Floating-Point Exceptions

Overflow, Underflow, Invalid, Precision, Denormal.

Other Exceptions

See Type E2; additionally

#UD If the EVEX broadcast bit is set to 1.
#UD If the MODRM.mod = 0b11.

VP4DPWSSD — Dot Product of Signed Words with Dword Accumulation (4-iterations)

Opcode/ Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEEX.DDS.512.F2.0F38.W0 52 /r VP4DPWSSD zmm1{k1}{z}, zmm2+3, m128	T1_4X	V/V	AVX512_4VNNIW	Multiply signed words from source register block indicated by zmm2 by signed words from m128 and accumulate resulting signed dwords in zmm1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
T1_4X	ModRM:reg (r, w)	EVEEX.vvvv	ModRM:r/m (r)	NA

Description

This instruction computes 4 sequential register source-block dot-products of two signed word operands with doubleword accumulation; see Figure 6-1 below. The memory operand is sequentially selected in each of the four steps.

In the above box, the notation of “+3” is used to denote that the instruction accesses 4 source registers based on that operand; sources are consecutive, start in a multiple of 4 boundary, and contain the encoded register operand. This instruction supports memory fault suppression. The entire memory operand is loaded if any bit of the lowest 16-bits of the mask is set to 1 or if a “no masking” encoding is used.

The tuple type T1_4X implies that four 32-bit elements (16 bytes) are referenced by the memory operation portion of this instruction.

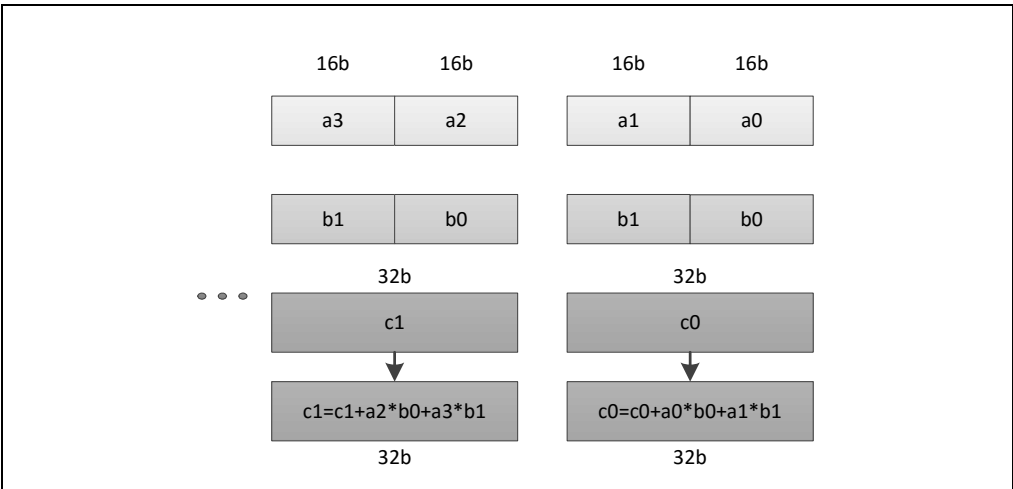


Figure 6-1. Register Source-Block Dot Product of Two Signed Word Operands with Doubleword Accumulation¹

NOTES:

1. For illustration purposes, one source-block dot product instance is shown out of the four.

Operation

src_reg_id is the 5 bit index of the vector register specified in the instruction as the src1 register.

VP4DPWSSD dest, src1, src2

(KL,VL) = (16,512)

$N \leftarrow 4$

ORIGDEST \leftarrow DEST

src_base \leftarrow src_reg_id & $\sim (N-1)$ // for src1 operand

FOR i \leftarrow 0 to KL-1:

IF k1[i] or *no writemask*:

FOR m \leftarrow 0 to N-1:

t \leftarrow SRC2.dword[m]

p1dword \leftarrow reg[src_base+m].word[2*i] * t.word[0]

p2dword \leftarrow reg[src_base+m].word[2*i+1] * t.word[1]

DEST.dword[i] \leftarrow DEST.dword[i] + p1dword + p2dword

ELSE IF *zeroing*:

DEST.dword[i] \leftarrow 0

ELSE

DEST.dword[i] \leftarrow ORIGDEST.dword[i]

DEST[MAX_VL-1:VL] \leftarrow 0

Intel C/C++ Compiler Intrinsic Equivalent

VP4DPWSSD __m512i _mm512_4dpwssd_epi32(__m512i, __m512ix4, __m128i *);

VP4DPWSSD __m512i _mm512_mask_4dpwssd_epi32(__m512i, __mmask16, __m512ix4, __m128i *);

VP4DPWSSD __m512i _mm512_maskz_4dpwssd_epi32(__mmask16, __m512i, __m512ix4, __m128i *);

SIMD Floating-Point Exceptions

None.

Other Exceptions

See Type E4; additionally

#UD If the EVEX broadcast bit is set to 1.

#UD If the MODRM.mod = 0b11.

VP4DPWSSDS — Dot Product of Signed Words with Dword Accumulation and Saturation (4-iterations)

Opcode/ Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEX.DDS.512.F2.0F38.W0 53 /r VP4DPWSSDS zmm1{k1}{z}, zmm2+3, m128	T1_4X	V/V	AVX512_4VNNIW	Multiply signed words from source register block indicated by zmm2 by signed words from m128 and accumulate the resulting dword results with signed saturation in zmm1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
T1_4X	ModRM:reg (r, w)	EVEX.vvvv	ModRM:r/m (r)	NA

Description

This instruction computes 4 sequential register source-block dot-products of two signed word operands with doubleword accumulation and signed saturation. The memory operand is sequentially selected in each of the four steps.

In the above box, the notation of “+3” is used to denote that the instruction accesses 4 source registers based on that operand; sources are consecutive, start in a multiple of 4 boundary, and contain the encoded register operand.

This instruction supports memory fault suppression. The entire memory operand is loaded if any bit of the lowest 16-bits of the mask is set to 1 or if a “no masking” encoding is used.

The tuple type T1_4X implies that four 32-bit elements (16 bytes) are referenced by the memory operation portion of this instruction.

Operation

src_reg_id is the 5 bit index of the vector register specified in the instruction as the src1 register.

VP4DPWSSDS dest, src1, src2

(KL,VL) = (16,512)

N ← 4

ORIGDEST ← DEST

src_base ← src_reg_id & ~ (N-1) // for src1 operand

FOR i ← 0 to KL-1:

IF k1[i] or *no writemask*:

FOR m ← 0 to N-1:

t ← SRC2.dword[m]

p1dword ← reg[src_base+m].word[2*i] * t.word[0]

p2dword ← reg[src_base+m].word[2*i+1] * t.word[1]

DEST.dword[i] ← SIGNED_DWORD_SATURATE(DEST.dword[i] + p1dword + p2dword)

ELSE IF *zeroing*:

DEST.dword[i] ← 0

ELSE

DEST.dword[i] ← ORIGDEST.dword[i]

DEST[MAX_VL-1:VL] ← 0

Intel C/C++ Compiler Intrinsic Equivalent

```
VP4DPWSSDS __m512i _mm512_4dpwssds_epi32(__m512i, __m512ix4, __m128i *);
VP4DPWSSDS __m512i _mm512_mask_4dpwssds_epi32(__m512i, __mmask16, __m512ix4, __m128i *);
VP4DPWSSDS __m512i _mm512_maskz_4dpwssds_epi32(__mmask16, __m512i, __m512ix4, __m128i *);
```

SIMD Floating-Point Exceptions

None.

Other Exceptions

See Type E4; additionally

- #UD If the EVEX broadcast bit is set to 1.
- #UD If the MODRM.mod = 0b11.

VPOPCNTD/VPOPCNTQ — Return the Count of Number of Bits Set to 1 in DWORD/QWORD

Opcode/ Instruction	Op/ En	64/32 bit Mode Support	CPUID Feature Flag	Description
EVEX.512.66.0F38.W0 55 /r VPOPCNTD zmm1{k1}{z}, zmm2/m512/m32bcst	FV	V/V	AVX512_VPOPCN TDQ	Counts the number of bits set to one in zmm2/m512/m32bcst and puts the result in zmm1 with writemask k1.
EVEX.512.66.0F38.W1 55 /r VPOPCNTQ zmm1{k1}{z}, zmm2/m512/m64bcst	FV	V/V	AVX512_VPOPCN TDQ	Counts the number of bits set to one in zmm2/m512/m64bcst and puts the result in zmm1 with writemask k1.

Instruction Operand Encoding

Op/En	Operand 1	Operand 2	Operand 3	Operand 4
FV	ModRM:reg (w)	ModRM:r/m (r)	NA	NA

Description

This instruction counts the number of bits set to one in each dword or qword element of its source (e.g., zmm2 or memory) and places the results in the destination register (zmm1). The EVEX encoded form of this instruction supports memory fault suppression.

Operation

VPOPCNTD

(KL, VL) = (4, 128), (8, 256), (16, 512)

FOR j ← 0 TO KL-1

IF MaskBit(j) OR *no writemask*

DEST.dword[j] ← POPCNT(SRC.dword[j])

ELSE IF *merging-masking*

DEST.dword[j] remains unchanged

ELSE

DEST.dword[j] ← 0

DEST[MAX_VL-1:VL] ← 0

VPOPCNTQ

(KL, VL) = (2, 128), (4, 256), (8, 512)

FOR j ← 0 TO KL-1

IF MaskBit(j) OR *no writemask*

DEST.qword[j] ← POPCNT(SRC.qword[j])

ELSE IF *merging-masking*

DEST.qword[j] remains unchanged

ELSE

DEST.qword[j] ← 0

DEST[MAX_VL-1:VL] ← 0

Intel C/C++ Compiler Intrinsic Equivalent

VPOPCNTD _mm<size>[_mask[z]]_popcnt_epi32

VPOPCNTQ _mm<size>[_mask[z]]_popcnt_epi64

SIMD Floating-Point Exceptions

None.

Other Exceptions

See Type E4.

INDEX

B

- Brand information 2-36
 - processor brand index 2-39
 - processor brand string 2-36

C

- Cache and TLB information 2-31
- Cache Inclusiveness 2-14
- CLFLUSH instruction
 - CPUID flag 2-30
- CMOVcc flag 2-30
- CMOVcc instructions
 - CPUID flag 2-30
- CMPXCHG16B instruction
 - CPUID bit 2-28
- CMPXCHG8B instruction
 - CPUID flag 2-30
- CPUID instruction 2-12, 2-30
 - 36-bit page size extension 2-30
 - APIC on-chip 2-30
 - basic CPUID information 2-13
 - cache and TLB characteristics 2-13, 2-31
 - CLFLUSH flag 2-30
 - CLFLUSH instruction cache line size 2-27
 - CMPXCHG16B flag 2-28
 - CMPXCHG8B flag 2-30
 - CPL qualified debug store 2-28
 - debug extensions, CR4.DE 2-29
 - debug store supported 2-30
 - deterministic cache parameters leaf 2-13, 2-15, 2-16, 2-17, 2-18, 2-19, 2-20, 2-21
 - extended function information 2-23
 - feature information 2-29
 - FPU on-chip 2-29
 - FSAVE flag 2-31
 - FXRSTOR flag 2-31
 - IA-32e mode available 2-23
 - input limits for EAX 2-24
 - L1 Context ID 2-28
 - local APIC physical ID 2-27
 - machine check architecture 2-30
 - machine check exception 2-30
 - memory type range registers 2-30
 - MONITOR feature information 2-34
 - MONITOR/MWAIT flag 2-28
 - MONITOR/MWAIT leaf 2-14, 2-15, 2-16, 2-17, 2-21
 - MWAIT feature information 2-34
 - page attribute table 2-30
 - page size extension 2-30
 - performance monitoring features 2-34
 - physical address bits 2-24
 - physical address extension 2-30
 - power management 2-34, 2-35, 2-36
 - processor brand index 2-27, 2-36
 - processor brand string 2-23, 2-36
 - processor serial number 2-30
 - processor type field 2-26
 - RDMSR flag 2-30
 - returned in EBX 2-26
 - returned in ECX & EDX 2-27
 - self snoop 2-31
 - SpeedStep technology 2-28
 - SS2 extensions flag 2-31

- SSE extensions flag 2-31
- SSE3 extensions flag 2-28
- SSSE3 extensions flag 2-28
- SYSENTER flag 2-30
- SYSEXIT flag 2-30
- thermal management 2-34, 2-35, 2-36
- thermal monitor 2-28, 2-31
- time stamp counter 2-30
 - using CPUID 2-12
- vendor ID string 2-24
- version information 2-13, 2-34
- virtual 8086 Mode flag 2-29
- virtual address bits 2-24
- WRMSR flag 2-30

E

EVEX.R 5-4

F

Feature information, processor 2-12

FXRSTOR instruction

- CPUID flag 2-31

FXSAVE instruction

- CPUID flag 2-31

I

IA-32e mode

- CPUID flag 2-23

L

L1 Context ID 2-28

M

Machine check architecture

- CPUID flag 2-30
- description 2-30

MAXSD- Return Maximum Scalar Double-Precision Floating-Point Value 1-3

MMX instructions

- CPUID flag for technology 2-31

Model & family information 2-34

MONITOR instruction

- CPUID flag 2-28
- feature data 2-34

MWAIT instruction

- CPUID flag 2-28
- feature data 2-34

P

Pending break enable 2-31

Performance-monitoring counters

- CPUID inquiry for 2-34

R

RDMSR instruction

- CPUID flag 2-30

S

Self Snoop 2-31

SpeedStep technology 2-28

SSE extensions

- CPUID flag 2-31

SSE2 extensions

- CPUID flag 2-31
- SSE3
 - CPUID flag 2-28
- SSE3 extensions
 - CPUID flag 2-28
- SSSE3 extensions
 - CPUID flag 2-28
- Stepping information 2-34
- SYSENTER instruction
 - CPUID flag 2-30
- SYSEXIT instruction
 - CPUID flag 2-30

T

- Thermal Monitor
 - CPUID flag 2-31
- Thermal Monitor 2 2-28
 - CPUID flag 2-28
- Time Stamp Counter 2-30

V

- Version information, processor 2-12
- VEX 5-1
- VEX.B 5-2
- VEX.L 5-2, 5-3
- VEX.mmmmm 5-2
- VEX.pp 5-2, 5-3
- VEX.R 5-3
- VEX.vvvv 5-2
- VEX.W 5-2
- VEX.X 5-2
- VPERMI2B - Full Permute of Bytes from Two Tables Overwriting the Index 6-1, 6-3
- VPERMT2B- Full Permute of Bytes from Two Tables Overwriting a Table 5-11
- VPERMT2W/D/Q/PS/PD—Full Permute from Two Tables Overwriting one Table 5-13
- VPMADD52HUQ—Packed Multiply of Unsigned 52-bit Unsigned Integers and Add High 52-bit Products to 64-bit Accumulators 5-20
- VPMADD52LUQ—Packed Multiply of Unsigned 52-bit Integers and Add the Low 52-bit Products to Qword Accumulators 5-18
- VPMULTISHIFTQB – Select Packed Unaligned Bytes from Quadword Source 5-22

W

- WBINVD/INVD bit 2-14
- WRMSR instruction
 - CPUID flag 2-30

X

- XFEATURE_ENALBED_MASK 2-1
- XRSTOR 1-1, 2-1, 2-35, 5-6
- XSAVE 1-1, 2-1, 2-4, 2-28, 2-35, 5-6